

Machine Learning for Risk-Aware Decision-Making in Short-Term Electricity Markets

On the Provision of Real-Time Balancing
Services in European Markets

by

Jérémie Bottieau

A thesis presented for the degree of
Philosophiae Doctor in Engineering Sciences & Technology

pursued in the Electrical Power Engineering Unit of the
University of Mons

Members of the Jury

Pr Marc Wuilpart, University of Mons Chairman

Pr Jalal Kazempour, Technical University of Denmark

Pr Fei Teng, Imperial College London

Dr Pieter Vingerhoets, Flemish Institute for Technological
Research (VITO)

Pr Thomas Brihaye, Université of Mons

Pr Zacharie De Grève, University of Mons

Pr Olivier Deblecker, University of Mons Secretary

Pr François Vallée, University of Mons Supervisor

Dr Jean-François Toubreau, University of Mons Co-supervisor

April 2022

Ce que l'on conçoit bien s'énonce clairement
Et les mots pour le dire arrivent aisément

extrait de l'Art poétique (1674) de Nicolas Boileau

Acknowledgements

Concrétiser cette thèse a été un long chemin, qui a été parsemé de doutes, de remises en question mais aussi de joies et de prises d'assurance. Ce chemin a débuté en vacances, plus précisément sur la place Saint Pierre du Vatican, où ma compagne, Zoé Vercruyssen, et moi avons eu la chic idée de suivre à la lettre les recommandations vestimentaires du Vatican sous une chaleur de 40°C. C'est à ce moment que nous avons rencontré d'une manière fortuite mon Prof. et futur promoteur, François Vallée, qui a réussi à me convaincre, certainement aidé par la chaleur, de l'immense opportunité que représente une thèse. Notez que, bien que le Pape s'apprêtait à donner l'Angélus, cet individu avait omis de me préciser les côtés les moins plaisants d'une thèse... Et donc, me voilà, quelques années plus tard, quelques papiers scientifiques écrits et quelques cheveux en moins, à me remémorer ce souvenir avec un sourire aux lèvres, heureux du chemin parcouru. Ainsi, à travers ces quelques mots, je voudrais remercier toutes les personnes qui m'y ont suivies, accompagnées ou rencontrées. Je me permets de commencer par des remerciements personnels avant d'aller vers des remerciements plus professionnels.

Tout d'abord, je tiens à remercier ma compagne, Zoé Vercruyssen (et je l'espère bientôt Bottieau), avec qui la vie n'est qu'une suite d'aventures rocambolesques plus ou moins contrôlées. Mère exemplaire de nos deux enfants (Louis et Emma), elle n'a eu de cesse de mettre notre famille au centre de nos préoccupations, mettant à leur juste place les soucis d'un thésard. Ton intransigeance, mais aussi ta douceur et le don de toi pour notre famille, ont été des éléments importants conduisant à notre bonheur actuel.

Ensuite, je remercie plus largement ma famille, notamment papy, papa, maman, mes frères, tata, les cousines et j'en passe, avec qui les bons comme les mauvais moments seront précieusement gravés dans ma mémoire. J'espère que nous pourrons garder aussi longtemps que possible cette force et cette envie d'être ensemble, qui nous font passer facilement au-dessus des épreuves de la vie. J'ai une pensée toute particulière pour mon papy qui, du haut de ses 90

ans, n'aura de cesse de m'impressionner par sa force d'esprit et son dos toujours bien droit.

Et puis, comment ne pas parler des collègues et amis peuplant le service de génie électrique, tous animés d'une flamme bien personnelle et atypique. J'ai bien aimé voir le service prendre de l'ampleur au fur et à mesure des années, et se transformer petit à petit en club professionnel de ping pong sous l'impulsion d'Ahmad. J'ai une note particulière pour Adriano Arrigo, ami de bureau, qui explique et clarifie instantanément des concepts mathématiques totalement abscons pour moi.

J'ai aussi une pensée pour le Prof. Kenneth Bruninx et le Dr Anibal Sanjab, avec qui j'ai eu le plaisir de collaborer grâce au projet EPOC 2030-2050. Interagir avec vous m'a certainement permis de gagner en confiance et de passer un cap dans ma recherche. Dans la même veine, je remercie le Prof. Yi Wang, avec qui j'ai eu la chance de collaborer en fin de thèse par le biais de mon co-promoteur, Dr Jean-François Toubeau. Les nouvelles thématiques de recherche qui y ont été abordées m'ont permis de finir cette thèse comme je l'ai commencée, en ayant faim d'apprendre.

D'une manière plus générale, je suis content que cette thèse m'a permis de contribuer au projet EPOC 2030-2050, qui est financé par le Fonds de transition énergétique organisé par le SPF Economie, PME, Classes Moyennes et Energie. J'espère que la réalisation de ce projet aboutira à des pistes concrètes pour la Belgique afin de l'accompagner au mieux vers une transition énergétique acceptable en termes de coûts.

Je remercie également les membres de mon Jury de thèse pour leur suivi, leur rigueur et la pertinence de leurs remarques scientifiques donnés tout au long de ma thèse. Chaque réunion de suivi a été une source d'apprentissage qui me poussait à progresser. Plus particulièrement, un grand merci aux intervenants extérieurs, Prof. Jalal Kazempour, Prof. Fei Teng et Dr Pieter Vingerhoets, pour votre retour scientifique hautement qualitatif.

Merci au Prof. Olivier Deblecker pour son exemplarité en tant que chef de service, assurant une certaine continuité et sérénité dans le service de Génie électrique. Je voudrais aussi spécialement remercier le Prof. Zacharie De Grève, qui a été un pilier important dans la réussite de cette thèse. Ayant eu la chance de t'accompagner sur divers projets industriels, j'ai beaucoup appris sur la gestion de projets et la communication de résultats scientifiques.

Enfin, comment puis-je terminer cette page de remerciements sans mentionner mes promoteur et co-promoteur, François Vallée et Jean-François Toubreau? Vous avez été des lanternes brillantes et bienveillantes dans la nuit que peut représenter la recherche. Au-delà de vos conseils, implication et amitié, je voudrais mettre particulièrement en avant votre attitude et votre travail au quotidien, qui ont été les sources principales d'inspiration dans la manière dont j'ai abordé la recherche. MERCI.

Summary

The European Union aims at reaching climate neutrality at horizon 2050 by achieving an economy with net-zero greenhouse gas emissions. Such ambitious policy target has implied and will continue to imply big changes in European energy systems, which accounted for 74% of European greenhouse gas emission in 2019. In this line, the electricity sector is currently rapidly shifting its supply mix towards less carbon-intensive energy sources, which conducts to growing shares of weather-dependent renewable energy sources (e.g., wind and photovoltaic power). The generation profile of these technologies differs from conventional ones (e.g., gas-fired units) by their high intermittency and uncertainty, which rises the difficulty of ensuring a continued, real-time balance between the electricity demand and supply.

This balancing requirement is vital when operating electricity systems since a mismatch between demand and supply automatically deteriorates the system frequency, which may trigger the disconnection of system components, and ultimately, lead to power blackouts. In the unbundled European electricity markets, the balancing management is supported via the balancing markets, which establish (amongst other) the market rules for the real-time trading of energy. While originally designed at national level, the European balancing markets are currently undergoing a harmonization process for fostering the cooperation between European countries. In this process, the favored option for pricing the real-time energy imbalances is the single price imbalance settlement, which provides financial incentives for market actors to adopt an imbalance position in the opposite direction of the total net system imbalance. When appropriately provided, this real-time balancing service is beneficial for the whole system as it reduces the total net system imbalance, which requires thus less costly corrective balancing actions.

This work is focused on developing novel forecast-driven strategies for fostering the provision of such real-time balancing services in European electricity markets. Practically, these strategies are studied using an integrated approach,

where the entire value chain, i.e., from forecasting to the decision-making processes, is modeled for optimizing close-to-real-time the imbalance position of a market actor. In this setting, the methodological contributions principally concern the modeling of uncertainty and risk in their operational strategies. More specifically, novel probabilistic forecasting methods based on deep learning techniques have been proposed, aiming at better capturing the high volatility of the total net system imbalance. In complement, for exploiting at best the predicted probabilistic information, tailored stochastic decision-support tools (i.e., stochastic programming and robust optimization method) are developed, for which a new data-driven approach was designed for continuously adjusting the risk policy of the market actor. The implementation of the developed approaches on real-world market data from the Belgian power system corroborates the key goal of the single price imbalance settlement, by showing that the market actor can increase its operating profit by optimizing its imbalance position, while reducing the total net system imbalance. Additionally, advanced neural networks architectures based on attention mechanisms demonstrate top forecasting performance for predicting the total net system imbalance. Finally, the data-driven approach for continuously adjusting the risk policy shows promising economic benefits in comparison with a static (determined once and for all) risk policy.

The final research efforts of this work are devoted on interpretability of deep learning-based forecasting methods, which aim at accurately identifying the most important input features of the model and their interaction when returning the prediction. Combining the predictive power of deep neural models with interpretable outcomes is an essential step for fostering their practical adoption in the energy industry. To achieve interpretability in both feature and time dimensions, the attention-based neural architecture is here augmented with subnetworks dedicated to endogenously quantify the relative importance of each input feature. Outcomes using the data from the Belgian power systems show that adding interpretable components within the neural architecture does not hinder their prediction performance, while shedding light on its most important drivers.

Contents

1. Introduction	1
1.1. Background	1
1.2. Motivation, Research Questions and Scope	5
1.3. Research Objectives and Scientific Contributions	10
1.4. Thesis Outline	13
2. European Short-Term Electricity Markets	15
2.1. Liberalization of the European Electricity System	17
2.2. System Operators, Regulators and Power Exchanges	20
2.3. Balance Responsible Parties	22
2.4. Balancing Service Providers	26
2.5. Clearing the Market	28
2.6. Towards an Integrated European Electricity Market	33
2.7. Trading in European Short-Term Electricity Markets	36
2.8. Provision of Real-Time Balancing Services	42
2.9. Conclusion	45
3. Towards High-Quality Probabilistic Time Series Forecasting Using Neural Networks	47
3.1. Time Series Forecasting	48
3.2. Forecasting System Imbalances	51
3.3. Feed-forward Neural Networks	53
3.4. Recurrent Neural Networks	57
3.5. Bridging Past and Future Input Variables	60
3.6. Training and Inference of Neural Networks	63
3.7. Benchmark	67
3.8. Metrics for Probabilistic Forecasting	69
3.9. Case study: Probabilistic Forecasting of Belgian System Imbalances	70
3.10. Conclusion	77

4. Risk-Aware Stochastic Provision of Real-Time Balancing Services	79
4.1. Market Assumptions and Application	80
4.2. The Bi-level Model	82
4.3. From the Bi-level Model to a Mixed-Integer Linear Formulation	86
4.4. Uncertainty Characterization	89
4.5. The Robust Optimization Model	90
4.6. The Stochastic Programming Optimization Model	92
4.7. Economic Interest of Accurate Probabilistic Forecasts	94
4.8. Economic Interest of the Bi-level Structure	99
4.9. Economic Interest of Risk Awareness for Providing the Real-Time Balancing Services	100
4.10. Conclusion	101
5. Automatic Risk-Adjusted Provision of Real-Time Balancing Services	103
5.1. Risk Awareness in Electricity markets Trading Strategies	104
5.2. Automatic Risk-Adjusted Decision-Support Framework	105
5.2.1. Goal of the Automatic Risk Adjustment Tool	106
5.2.2. Training and Inference Stages of the Risk Adjustment Tool	107
5.2.3. Machine-Learning Models	108
5.3. Case Study	112
5.4. Conclusion	120
6. Towards Interpretable Probabilistic Forecasting Using Neural Networks	123
6.1. Interpretability in Neural Models	124
6.2. Probabilistic Forecasting of Imbalance Prices	126
6.3. Interpretable Transformer Model	128
6.3.1. Feed-Forward Networks	130
6.3.2. Variable Selection Layer	131
6.3.3. Local Temporal Processing Layer	132
6.3.4. Transformer-Specific Attention Layer	133
6.3.5. Regularization layer	134
6.3.6. Output layer	136
6.4. Benchmark	136
6.5. Case Studies	138
6.5.1. Forecast Evaluation	138
6.5.2. Ablation Analysis	141
6.5.3. Interpretability	143
6.6. Conclusion	147

7. Conclusions and Perspectives	149
7.1. Conclusions	149
7.2. Perspectives	153
A. List of Publications	175

CHAPTER 1.

Introduction

The first chapter provides a brief introduction to the current energy challenge, which allows positioning the objectives and contributions of the proposed research work. More specifically, Section 1.1 gives general information on the importance of energy in human activities and its effect on economy and environment, for contextualising the 2050 European Union’s energy development goals, which emphasises the key role of electricity as a decarbonized energy carrier. On this basis, Section 1.2 presents the motivation, research questions and scope of the proposed research work. Then, Section 1.3 exposes the associated research objectives and scientific contributions. Finally, Section 1.4 provides an outline of the remainder of this dissertation.

1.1. Background

Energy is a key competency for the economic development of a country, whose abundance greatly facilitates the daily activities of human life. Before the 1900s, the daily human activities were essentially powered by the burning of biomass (mostly wood), the mechanical force deployed by small-sized windmills, or simply by human and animal muscles. Then, at the turn of the 20th century, novel energy sources based on the combustion of fossil fuels (i.e., coal, oil and natural gas) were progressively discovered and introduced in the daily human activities [1]. This is illustrated in Fig. 1.1a, which provides a worldwide picture of the use of energy over timepoints of the last two hundred years. It can be observed that the use of fossil-based energy have steadily increased, and that such an increase was also accompanied with a rapid economic growth at world level (which is showcased in Fig. 1.1b). Fig. 1.1b estimates the worldwide economic growth using the per capita gross domestic product (GDP), i.e., the total monetary income per person of the goods and services produced within a defined geographical area, which is evaluated in constant 2011 \$ over the years




Figure 1.1.: World energy use (Fig. 1.1a) and economic growth (Fig. 1.1b) over the last two centuries [2].

1820-2018. During this time period, the GDP indicator rose by more than a factor of twelve at world level, reaching 15212 \$ in 2018 from 1102 \$ in 1820. Besides, this increase of GDP is even more spectacular for Western Europe and the United States, where their GDP stepped up by a factor of about twenty [2].

During this relatively short period of time, the world entered in a new economic era, which was supported by constant innovations (e.g., the steam engine, electricity or combustion engine) and the intensive use of fossil fuels. From a positive perspective, the resulting economic growth have enabled important social progresses (e.g., on literacy, infant mortality, and life expectation), an increased availability of products, and a facilitated mobility, which allows granting the majority of the population higher living standards. However, the deposits of fossil fuels are physically bounded and are not easily replenished, while their intensive exploitation and massive combustion negatively impact the environment. Indeed, scientific evidences attribute the greenhouse gases (e.g., carbon dioxide CO_2) emitted by the combustion of fossil fuels as key drivers for the climate change, which tends to a relatively rapid global warming of the Earth (see Fig. 1.2). A consensus position for avoiding the worst consequences of global warming is to limit the average temperature rise to less than 2°C compared to pre-industrial levels [3]. This requires a rapid shift from a fossil-based economy towards a carbon neutral society, which unleashes less greenhouse gas emissions in the atmosphere. Such considerations have increased the shares of nuclear, hydropower and renewable energies in the energy supply mix over the past half century, but their input at world level is still limited at 18.5% in 2021 according to [4].

As about 74% of the greenhouse gas emissions globally originate from the

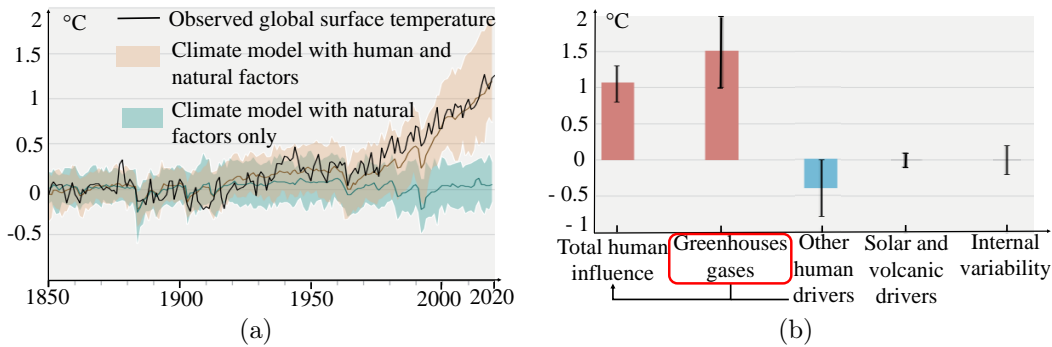


Figure 1.2.: Changes in global surface temperature over the past 170 years (Fig 1.2a) and aggregated contributions to 2010–2019 warming relative to 1850–1900, assessed from attribution studies (Fig. 1.2b) [5].

energy sector in Europe (see reference [4] over the year 2019), the European Union has drawn up a road map with joint energy and climate goals to be reached for the years 2030 and 2050, which are currently revised by the European Commission in the ‘Fit for 55’ package [6]. At each step, quotas in terms of reduction in greenhouse gas emissions, development of renewable energies and improvement in energy efficiency are defined. Hence, by 2030, the Member States of the European Union will be enforced to: i) reduce their greenhouse gas emissions by 55% compared to 1990 levels, ii) increase the energy efficiency of their economies by 36-37% compared to a 2007 baseline, and iii) increase the share of renewable energies in their final energy consumption to 40%. In 2050, the European target is to achieve a climate-neutral economy, i.e., an economy with net-zero greenhouse gas emissions. The achievement of these climate and energy goals constitutes a major challenge for the energy sector, which comes down to effectively shift fossil-based energy sources towards more carbon-free energy sources (e.g., the renewable or nuclear energy sources), while ensuring a security of supply and competitive prices for the end-user.

In this context, electricity is considered as a crucial energy carrier for reaching the carbon neutral economy target in 2050. Formerly, electricity was favoured for its intrinsic qualities such as high final conversion efficiencies, flexibility in its development planning, precise control of the output power, and property that no mechanical noise is emitted at the point of use. Today, this trend is further strengthened by the possibility of incorporating less carbon-intensive technologies in the electricity supply mix, with, for instance, wind and photovoltaic production. In complement, an electrification process at the demand-side is currently boosted in the European Union, where novel electricity-based solu-




Figure 1.3.: European electricity generation for renewable energies over 1990-2019 [4]

tions in the transportation and residential heating sectors (e.g., electric cars or electric heat pumps) are currently strongly encouraged [7].

Fig. 1.3 illustrates the shift initiated in the electricity supply mix, where the European electricity production from renewable energy sources is continuously rising since the 1990s. However, this massive introduction of Variable Renewable Energy Sources (VRES)¹, like wind and photovoltaic production, greatly impacts the planning and operation of the electricity systems [8], [9]. Practically, a stable and secure operation of an electricity system requires a continued balance in real-time between the electricity demand and supply. The difficulty of managing this balance requirement is rising in today's electricity systems with the growing shares of VRES. Indeed, as their production profiles are highly intermittent and uncertain (due to the chaotic atmospheric conditions governing them), the VRES cannot be straightforwardly piloted for following the variations in electricity demand. In addition, the VRES are characterized by high capital and low operating costs, whose investment costs are largely predominant over fuel (which is free) and maintenance. Hence, the massive introduction of VRES in electricity systems was initially triggered by support mechanisms, which have distorted the wholesale electricity market prices due to their operational costs close to zero. These lower prices tend to push off-market more rapidly than anticipated conventional power plants (e.g., gas-fired units), even though the electricity systems are still designed to be operated with such controllable units. This brings major and pressing

¹also denoted as weather-dependent renewable energy sources

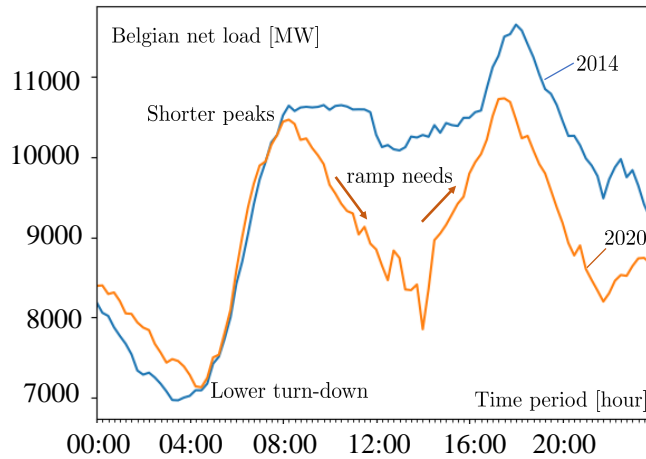


Figure 1.4.: The Belgian net load for the 6th January 2014 and 2020.

challenges in the planning and operation of electricity systems, which now requires to manage the variability and limited predictability of VRES, while less controllable production units are available [10], [11].

1.2. Motivation, Research Questions and Scope

For ensuring a well-functioning electricity system, two critical aspects at different time horizons are typically considered: adequacy and security [12]. Adequacy indicates the ability of the electricity system to own a sufficient production capacity (including imports) to supply with high enough probability the end-users at all times. This aspect principally explores whether enough infrastructure investments are initiated in the electricity system with a long term perspective (e.g., a 10-year period) [8]. Security refers to the ability of the system to respond to sudden disturbances arising in the electricity system without major service interruptions [13]. This aspect principally explores whether the electricity system can recover from sudden, unanticipated events such as electric short circuits or non-anticipated loss of system components with a short term perspective. In this work, the interest is on the short-term operation of the electricity system, i.e., the ability of the electricity system to modify its electricity production or consumption in response to variability (expected or not). Short-term operational issues have grown in significance lately since these issues have been intensified with the increasing penetration of weather-dependent renewable energy sources.

Short-term operational issues were formerly driven by i) outages on the supply side and the transmission networks, and ii) the variable and uncertain electricity demand. Forced outages, which results in a sudden loss (or excess) of power, are an inherent characteristic of generation and transmission systems, and are hardly predictable. Then, the exact electricity demand is difficult to know beforehand as it depends on external variables such as consumer preferences and weather conditions. Today's electricity systems face an additional source of variability and uncertainty, which is introduced by the weather-dependent renewable energy sources. The net load curve can be used for illustrating one aspect of this variability. The net load consists in subtracting the wind and photovoltaic electricity production from the total load of the electricity system. In certain times of the year, with the increased contribution of weather-dependent renewable production, these curves produce a 'belly' appearance in the mid-afternoon that quickly ramps up to produce an 'arch' similar to the neck of a duck. This type of curve is called the 'duck curve' [14]. In Fig. 1.4, this curve is produced for the Belgian power system on a quarter hourly basis for the 6th January of both 2014 and 2020 years, which are respectively characterized by 4402 MW and 7682 MW of weather-dependent renewable production capacities. We can observe that, for that day, the system operator needs to be more reactive in 2020 with the controllable electricity production for meeting the sharp changes in the electricity net load. Two distinct ramp periods emerge. The first one in the downward direction, occurring after 8:00 a.m., when conventional generation is replaced by photovoltaic production (producing the belly of the duck). The second ramp occurs around 5:00 p.m. when the sun lies down and photovoltaic production ends, which require the system operator to dispatch rapidly resources in the upward direction (the arch of the duck's neck). Hence, the introduction of weather-dependent renewable generation reduces the number of full load hours of controllable units, which now requires novel high flexible operational capabilities for starting and stopping multiple times per day. Such effects substantially rise the difficulty of ensuring a continued, real-time balance between the electricity demand and supply.

However, this balancing requirement is vital when operating electricity systems as a mismatch between demand and supply automatically deteriorates the system frequency (the nominal frequency is set at 50 Hz in Europe), which may consequently trigger the disconnection of system components, and ultimately, lead to power blackouts. In order to maintain the system in balance, the expected and unexpected variations of net load have to be covered with flexible means of the electricity system, which can alter their production or demand upon a signal request in a relatively short-time frame. As showcased in Fig. 1.5, this can be provided by:

- controllable thermal power plants: the traditional source of short-term operational flexibility in power systems is delivered by the cycling of controllable thermal power plants (e.g., coal- and gas-fired steam power plants). Cycling consists in starting up, shutting down, ramping up or ramping down a power plant to vary its power output [9]. Of course, such a dynamic operation is accelerating the unit's aging, which rises its operational cost [15];
- demand response: demand response accommodates the demand of an electricity load following a signal (e.g., a price signal). Natural candidates for demand response are i) thermostatically controlled loads, which can contain a certain thermal inertia. This allows these loads to be both fully responsive and non-disruptive in terms of the perceived energy service [16]. ii) Electric vehicles, whose charging profiles of their batteries can be managed when parked [17]. Although industrial demand response can be currently exploited in electricity systems, the best pathway for introducing residential demand response is still subject to researches [18], [19];
- energy storage systems: these technologies are characterised by an energy reservoir (which can store the energy via another carrier), which bidirectionally exchanges with the electricity system. Different energy storage technologies exist [20], where pumped-hydro storage units and battery facilities are the most well-developed [21], [22];
- interconnections: interconnections allow to import (or export) flexibility from/to other regions by means of cross-border exchanges. The development of interconnections goes hand in hand with an integrated European electricity market, as it allows to facilitate the energy exchanges between European countries [23].

Note that the weather-dependent renewable energy sources could also be part of the solution by relying on adequate control techniques of their power electronics devices located at the interface between the installation and the electricity network. Hence, for instance, wind turbines under specific control schemes – see, e.g., in [24] – can modulate their output power, and even deliver balancing services to the grid if they are operated below their maximum output power (although, the reliability of the committed service may be questioned due to the inherent intra-hour wind speed fluctuation) [25].

Thermal power plants, demand response, energy storage systems, interconnections and active control of weather-dependent renewable energy sources are currently the main flexibility options in today's electricity systems. Since the uncertainty of weather-dependent renewable energy sources decreases closer to the moment of delivery [26], [27], adjusting the decisions of flexible power units

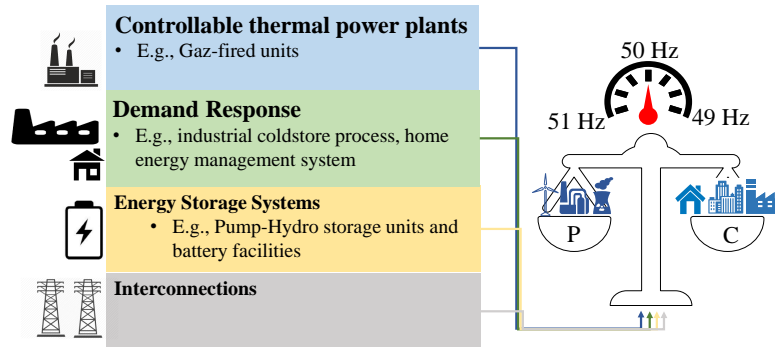


Figure 1.5.: The principal flexible means of the electricity system.

closer-to-real-time is needed for unlocking the cost-saving potential associated with this gain of accuracy – see, e.g., [28], [29]. Competing in liberalized electricity markets (see Chapter 2), operators of power units, which are thus market actors, currently adapt their short-term physical operations based on the financial exchanges performed on the electricity markets [30]. Hence, the decisions made on the electricity markets at a short-(i.e., one day in advance) and very-short-(i.e., few hours to a quarter hour in advance) terms are thereby crucial in order to appropriately manage a flexible portfolio (e.g., batteries, pumped-storage stations, etc.). Yet, as very-short-term electricity markets (which are all more detailed in Chapter 2) suffer from a lack of liquidity, they are not pragmatically considered in the literature and, thus, remain insufficiently researched to date [31], [32]. This status is currently moving as these markets have recently received a special attention from European policy makers to increase their harmonization and, consequently, their liquidity. For instance, different market pilots are currently launched in Europe such as the cross-border intraday market (XBID) project [33] or various initiatives concerning the balancing markets (e.g., PICASSO or TERRE projects) [34].

The balancing markets are the main mechanisms for supporting the real-time balancing requirement in the European electricity systems. While originally designed at national level, the rules of the balancing markets follow an harmonization process for fostering the balancing cooperation between European countries [35]. In this process, the favored option for pricing the real-time energy imbalances is the single price imbalance settlement, which provides financial incentives to market actors to adopt a real-time out-of-balance position in the opposite direction of the total net system imbalance. These real-time balancing services, when appropriately provided, allows reducing the total net system

imbalance, and is thus beneficial for the whole system, which requires less costly corrective balancing actions. In this sense, Brijs et al. [36] attribute the increased provision of such real-time balancing services as one of the reasons of the decline of balancing requirements in the Central Western European region from 2012 to 2015. In this context, several authors have studied near real-time dispatch strategies for providing these real-time balancing services in a single price imbalance setting [37]–[39]. While these strategies are presented in more details in Section 2.8 of Chapter 2, their common feature is that they rely on deterministic decision-support tools, wherein optimal decisions, i.e., the intentional imbalance positions of the actor, are based on a single-point forecast of the system imbalance direction (which is positive or negative).

However, the real-time market signals such as the total net system imbalance have proven to be highly volatile and difficult to predict [40], [41], and sub-optimal out-of-balance positions can expose the actors to significant financial penalties [42]. The high uncertainty and risk associated with the provision of real-time balancing services create thus great challenges for the market actor. Besides, the volume of energy traded in real-time is by design small (as it describes a mismatch between the electricity demand and supply), such that the provision of even a small real-time balancing service (in terms of volume of energy) may impact the real-time system balancing. For all these reasons, this thesis aims at developing novel formulations for answering the main following research question:

- *How can we efficiently support the decision-making process of a flexible actor willing to provide real-time balancing services to the electricity system via the opportunities provided by the single imbalance pricing scheme?*

Throughout this work, this main research question has been subdivided into four subquestions, and consequently transferred into the research strategy. Each of the following subquestions is addressed in the following chapters of this dissertation:

- *How can we improve the forecasting performance of future system imbalances, with an accurate quantification of the level of uncertainty?*
- *How do we mathematically formulate the optimal decision-making process of a strategic actor providing real-time balancing services, while taking into account both the market small size and high uncertainty?*
- *Being exposed to significant financial penalties, can we improve the financial risk management of a market actor providing real-time balancing services?*

- *Besides providing high-quality probabilistic predictions, can we provide interpretable insights on how they are generated for better increasing the designer and/or user confidence in the prediction outcomes?*

Overall, the generic nature of the associated methodological contributions pertaining to the modeling of uncertainty and risk in operational strategies of market actors may nourish further research efforts in other market frameworks. Besides, the achievement of this research project allows gaining new or improved insights on the provision of real-time balancing services within a single imbalance price scheme (which is the current European favored option). The case studies are based on real-world market data from the Belgian power system, which allows performing quantitative observations for this market segment. The presented research contributions and findings may thus be of interest to all electricity actors, e.g., researchers, policy makers, system operators, business entities.

1.3. Research Objectives and Scientific Contributions

The goal of this thesis is to develop novel forecast-driven strategies for fostering the provision of real-time balancing services. Practically, these strategies are studied using an integrated approach, where the entire value chain, i.e., from forecasting to the decision-making process, is modeled for optimizing close-to-real-time the imbalance position of a market actor. More specifically, the presented work sets the following four research objectives:

- **Attaining high-quality probabilistic forecasts.** High-quality probabilistic forecasts of the total net system imbalance are a mandatory condition for a well-informed intentional imbalance position of the market actor within the single price imbalance settlement framework. In this objective, novel deep learning-based time series forecasting tools are investigated for generating improved probabilistic predictions.
- **Developing tailored risk-aware stochastic decision-support tool.** Based on the predicted probabilistic information, novel mathematical formulations are developed for incorporating i) the small market size, and ii) the high uncertainty and risk associated with these real-time balancing opportunities.
- **Adjusting continuously the risk policy of a market actor.** In this objective, machine learning techniques are leveraged for adjusting, at each decision step, the financial risk management of an actor based on the current market conditions along with its expected objective outcomes.

- **Adding interpretability in probabilistic forecasts.** This objective aims at incorporating notions of interpretability (i.e., the identification of the most important input features over time) in deep learning-based time series forecasters for improving the designer and/or user confidence in their outcomes.

The achievement of these four complementary objectives provides an integrated decision-making tool for optimizing close-to-real-time the energy position of a market actor's portfolio in European electricity markets. From a practical point of view, this research project would enable market participants to mobilize their remaining margins of flexibility that do not meet technical requirements (e.g., in the balancing markets) and/or sufficient economic gains (e.g., in the day-ahead market) in previous market segments. By extension, this provision of additional flexibility to the system would also benefit the end-users by reducing the balancing activation costs. Given these research objectives, we list the four main contributions w.r.t. the literature:

- **On accuracy in probabilistic forecasting.** This thesis investigates and pushes further deep neural network architectures for predicting real-life time series of electricity systems (e.g., the total net system imbalance and real-time electricity prices). Particularly, advanced sequence-to-sequence recurrent neural models are developed aiming at better capturing the complex spatio-temporal dynamics characterizing electricity quantities. The proposed advancements have resulted in the following contributions:
 - **J. Bottieau**, F. Vallée, Z. De Grève, J-F. Toubreau, "Leveraging Provision of Frequency Regulation Services from Wind Generation by Improving Day-Ahead Predictions using LSTM Neural Networks," in 2018 IEEE International Energy Conference (ENERGYCON), Limassol, Cyprus, 2018.
 - J-F. Toubreau, **J. Bottieau**, F. Vallée and Z. De Grève, "Deep Learning-Based Multivariate Probabilistic Forecasting for Short-Term Scheduling in Power Markets, " in IEEE Trans. Power Syst., vol. 34, no. 2, pp. 1203-1215, 2019.
 - J-F. Toubreau, T. Morstyn, **J. Bottieau**, K. Zheng, D. Apostolopoulou, Z. De Grève, Y. Wang, and F. Vallée, "Capturing Spatio-Temporal Dependencies in the Probabilistic Forecasting of Distribution Locational Marginal Prices," in IEEE Trans. Smart Grid, vol. 12, no. 3, pp. 2663-2674, 2021.
- **On tailored risk-aware decision-support tools.** This thesis proposes two novel decision-support tools for providing real-time balancing services: i) one based on stochastic programming, which optimizes the imbalance

position of a market actor to maximize its expected profit over a finite set of scenarios, and ii) one based on robust optimization, which optimizes the imbalance position of a market actor under the worst-case realization of the probabilistic predicted information. Both formulations are structured as a bi-level program, which mathematically capture the interaction between the intentional imbalance position of the actor and the price signal of the single imbalance pricing mechanism. This part led to the following contribution:

- **J. Bottieau**, L. Hubert, Z. De Grève, F. Vallée and J-F. Toubeau, "Very-Short-Term Probabilistic Forecasting for a Risk-Aware Participation in the Single Price Imbalance Settlement, " in *IEEE Trans. Power Syst.*, vol. 35, no. 2, pp. 1218-1230, 2020.
- **On adjusting the risk policy of a market actor.** For the first time, a data-driven risk-adaptive approach is implemented for progressively updating and improving the risk policy of a market actor based on the dynamically changing market operating conditions. This approach is tested on the scenario-based and robust optimization frameworks. This methodology has been published in:
 - **J. Bottieau**, K. Bruninx, A. Sanjab, Z. De Grève, F. Vallée and J-F. Toubeau, "Automatic Risk Adjustment for Profit Maximization in Renewable Dominated Short-Term Electricity Markets," in *ITEES*, vol. 3, issue 12, 2021
- **On interpretability in probabilistic forecasting.** Novel attention-based neural architectures are proposed for predicting real-time electrical quantities in an interpretable fashion. The selection of relevant variables is internalized into the model, which provides insights on the relative importance of individual inputs. Then, attention mechanisms are used for enabling the model to explicitly focus on relevant temporal dependencies, which allows shedding light on the most relevant time dynamics such as seasonal patterns. In this line, two contributions have been proposed (one is published, and the other one is currently peer-reviewed):
 - **J. Bottieau**, Y. Wang, Z. De Grève, F. Vallée, and J-F. Toubeau, "Transformer Model for Interpretable Probabilistic Forecasting of Real-Time Electricity Prices," currently in second round of review in *IEEE Trans. Power Syst.*
 - J-F. Toubeau, **J. Bottieau**, Y. Wang, and F. Vallée, "Interpretable Probabilistic Forecasting of Imbalances in Renewable-Dominated Electricity Systems," in *IEEE Trans. Sust. Energy*. Early access: <https://ieeexplore.ieee.org/abstract/document/9464660>.

The complete list of publications (including side projects) can be found in Appendix A.

1.4. Thesis Outline

The organisation of this report is presented as follows:

- **Chapter 2** firstly introduces the basics concerning the European short-term electricity markets, and is concluded by a detailed presentation of the opportunities offered by the provision of real-time balancing services.
- **Chapter 3** is focused on the probabilistic time series forecasting models, with an emphasis on neural networks models. The developed models are assessed and compared on the Belgian system imbalance signal.
- **Chapter 4** presents the risk-aware stochastic decision-support tools for providing the real-time balancing services.
- **Chapter 5** extends the risk-aware stochastic formulations by proposing an automatic risk-adaptive approach for continuously adjusting the financial risk management of a market actor.
- **Chapter 6** details a (neural) Transformer-based model that provides interpretable prediction outcomes. The forecasting model is tested on the Belgian (real-time) imbalance prices, which are typically characterized by low- and high-price regimes. The case study allows shedding light on the attribution importance of the different inputs processed by the model depending on the predicted price regime.
- **Chapter 7** concludes this report, while proposing some perspectives for future research.

CHAPTER 2.

European Short-Term Electricity Markets

Electricity is a commodity with the following characteristics:

- Electricity demand and supply need to be matched continuously for ensuring a stable frequency within the electricity network.
- Electricity follows paths defined by physical laws (e.g., Kirchhoff's laws), and not the pathways defined by financial transactions.
- Electricity flows through a complex infrastructure, wherein a wide range of components (e.g., alternators, transformers, protection devices, inverters, physical lines and industrial motors) interconnects the production units to the end-users across a wide geographical area.

Fig. 2.1 showcases the synchronous electricity network of Continental Europe, which serves over 400 million customers in 24 countries. The physical infrastructure of electricity systems is commonly subdivided into i) the high-voltage transmission system, which carries the electricity generated by large power plants over long distances (typically at 380kV - 30kV), and ii) distribution systems, which sources residential and industrial consumers at lower voltage levels (from 30kV to 220V). The continuity of the electricity service is vital to all of connected end-users for their daily activities (e.g., hospitals or cold storages). If a mismatch between electricity demand and supply appears, the network encounters difficulties such as deviations of the system frequency (set at 50 Hz in Europe, as illustrated in Fig. 2.1), which can disturb the optimal operation of the system components (e.g., industrial motors) and the grid assets (transformers, power plant auxiliary equipment, etc.). Indeed, in case of generation shortage, the missing electricity is taken from the inertia of the rotating machines that are synchronized with the grid. These generators are thus decelerating, which leads to a frequency decrease, and even

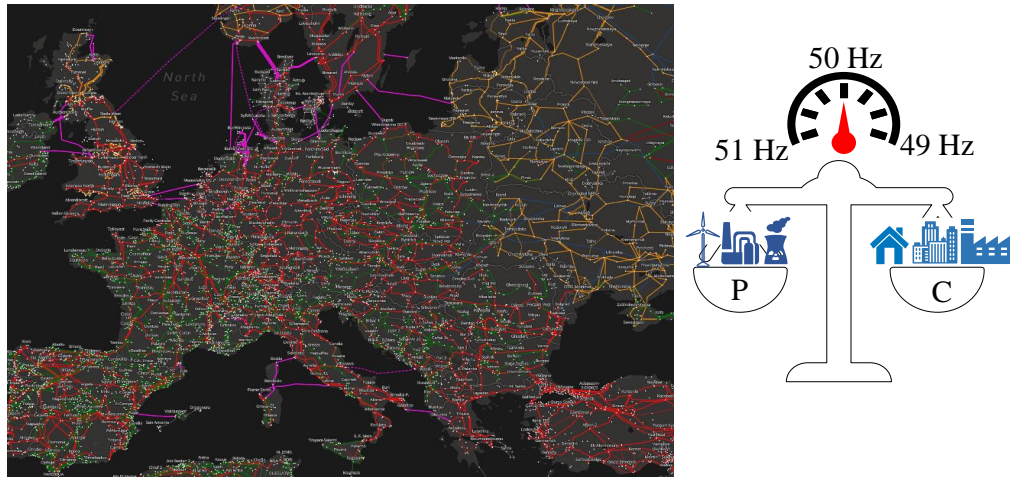


Figure 2.1.: The synchronous grid of Continental Europe. [44].

to a disconnection of these units if the rotation speed goes below a threshold. Consequently, if a frequency deviation is not immediately alleviated, the electricity network can face a domino effect of disconnections of generators, which importantly jeopardizes its stability. Ultimately, a failure to restore the balance could lead to an entire system collapse, which is called blackout. A blackout is very costly for society as a whole as it takes several hours for restoring the electricity services. For instance, it took about 12 hours to restore the electricity system after the Great Blackout of 2011 in Pacific Southwest affecting 2.7 million residents of San Diego, California, Arizona and Mexico [43].

Due to the characteristics of electricity, the planning and operation of electricity systems are not straightforward. This complexity is reverberated in the price of electricity, which includes, besides the price of commodity, other regulated costs related to the well-functioning of the network infrastructure (the so-called transmission and distribution costs). Moreover, additional taxes and levies, which promote specific energy development policies, can be added on top of the electricity price. **In this work, we solely focus on the commodity component of the electricity price, which is obtained via the wholesale electricity markets.**

The liberalization process of the European electricity system is discussed in Section 2.1. This liberalization process led to the introduction of market mechanisms for organizing the exchanges between self-interested third parties, where new roles and responsibilities have been defined. The roles of system operators, regulators and power exchanges are discussed in Section 2.2, while

the different responsibilities that are endorsed by market actors, i.e., Balance Responsible Parties and Balancing Service Providers, are respectively presented in Section 2.3 and 2.4. The electricity markets aim at maximizing the coordination between the electricity supply and demand sides in a competitive framework, while considering to the best extent the physics governing the power units and flows of electricity. Section 2.5 describes how both electricity supply and demand sides are accorded in this framework, while detailing how the physical constraints governing electricity systems are introduced in the markets. These electricity markets are now well-anchored in each Member State, but harmonization efforts are still required for defining a common structure across Europe. The current status of this harmonization process is provided in Section 2.6 for the European short-term electricity markets. Then, Section 2.7 presents an overview of the different challenges and trading solutions associated with each short-term electricity market, while Section 2.8 presents a more detailed discussion of the targeted application, i.e., the provision of real-time balancing services. Finally, Section 2.9 concludes this Chapter.

2.1. Liberalization of the European Electricity System

In most European countries, the electricity system was initially planned and operated by single, vertical national utilities. Each of these national entities was responsible for the generation, transmission, distribution, and retail of electricity for their geographical area. At that time, as the supply of electricity was considered as a service of general economic interest, it was not subject to the rules of competition established by the European Union (EU) Treaties of Rome (1957) and Maastricht (1993). In the 1990s, a wave of liberalization concerning different industrial sectors (e.g., air transport, insurances or banking) started in Europe, driven by the promise of efficiency earned via lower costs and prices. In line with these sectors, the electricity system was also reformed as follows: i) the generation, purchase and sale of electricity were unbundled from transmission and distribution services, ii) markets were organized for introducing competition between producers and suppliers, iii) regulated monopolies operating the transmission and distribution networks were established, and iv) regulators monitoring both the market-based and regulated activities were instituted [45]. Since the beginning of this liberalization process, the achievement of a common internal electricity market across all Member States was envisaged [46]. However, after thirty years of legislation, the completion of this internal electricity market is still not achieved, which reflects the difficulty of reforming a patchwork of electricity systems that were centrally planned. As shown in Fig. 2.2, the liberalisation process was driven by four European

1996	2003	2009	2019
1 st Energy Package <ul style="list-style-type: none"> • Functional unbundling • Wholesale electricity markets 	2 nd Energy Package <ul style="list-style-type: none"> • Functional and legal unbundling • Creation of National Regulation Authorities • Retail markets 	3 rd Energy Package <ul style="list-style-type: none"> • Definition of unbundling models • Creation of ACER • Creation of ENTSO-E 	4 th Energy Package <ul style="list-style-type: none"> • Updates of EU market rules • Energy communities • Launch of the EU DSO entity • Risk preparedness for future electricity crises

Figure 2.2.: The timeline of the different energy packages that support the liberalization process of the European electricity system [45].

legislative energy packages over the period 1996 to 2019. The energy packages were then transposed into national laws by each Member State in the years following. It should be noted that the same liberalization process has happened for the gas sector.

The First Package in 1996 takes the first steps towards the unbundling of the competitive part (producers and suppliers) from the regulated part (the transmission and distribution activities). The collusion between these two parts is hampered by prohibiting the transmission/distribution system operators, i.e., the regulated entities managing the high-voltage (380kV - 30kV) and medium to low-voltage (30kV - 220V) electricity networks, to disclose any confidential information to the competitive side, while being independent (at least functionally) from other third parties. However, a large room of manoeuvre is left for each Member State for introducing and organizing the wholesale competition between producers, large industrial consumers, and electricity retailers, which results in different national electricity markets.

The Second Package in 2003 amends the previous one by going a step further in unbundling requirements for the transmission and distribution system operators. Indeed, the electricity network should now be operated, maintained and developed through a legally separate entity. National Regulation Authorities are also designated for ensuring transparency and competitiveness in the electricity markets, while guaranteeing the public interest. Besides, smaller industrial and residential end-users were now free to choose their own electricity suppliers from a panel of actors.

The Third Package focuses on resolving certain structural problems, while improving the functioning of European electricity markets. Indeed, the final

report of the energy sector inquiry in 2007 identifies the lack of integration between national markets and their high degree of concentration as important persisting shortcomings [47]. In this line, for alleviating discrimination access to the grid that may obstruct competitors' entrance, this package further strengthened the unbundling process by presenting three models, i.e., ownership unbundling, independent system operator, and independent transmission system operator. The ownership model was openly favoured, which requires that the transmission system operator is completely independent from producers and suppliers. An Agency for the Cooperation of Energy Regulators (ACER) at European level is also established for increasing the cooperation between the different national regulation authorities. Similarly, the European Network for Transmission System Operators for Electricity (ENTSO-E) is also built for ensuring an optimal management between the different national transmission system operators. Through both entities, the Third Energy Package triggered the creation of electricity network guidelines and codes, which technically details common guidelines and rules for designing the European electricity markets.

Lastly, in the line of the 2015 Paris Agreement, a Fourth Package, i.e., the "Clean Energy for All European" package, was implemented in 2019. This package is more oriented toward the energy transition and climate goals, with a focus on the participation of citizens in the energy transition. Through eight legislative acts, the European Union updates its energy policy framework for reaching the targets of 2030, i.e., 40% cut in greenhouse gas emissions compared to 1990 levels, 32% of renewable energy sources in the energy mix and 32.5% of increase in energy efficiency over current levels. To do so, amongst others, this package reworks the European market rules for improving the level playing field, i.e., putting all technologies and market actors on equal footing. It also introduced the concept of energy communities in its legislation, which enables the local active participation of citizens in the electricity system. It also increases the competences of the European Union agency ACER for, e.g., drafting the network codes and guidelines. It also mandates the establishment of an European distribution system operator entity called EU DSO entity, and provides guidance regarding the establishment of capacity mechanisms, i.e., a mechanism aiming at ensuring adequate investment in power plants. Besides, it obliges each Member State to prepare plans for preventing and managing potential future electricity crises.

All this legislative process leads to an organisation of the economic dimension of the electricity systems into different markets, ranging from long-term to short-term arrangements. In the following Sections, the roles and responsibilities of the different actors are presented.

TSO :	Transmission System Operator	
NRA :	National Regulatory Authority	
DSO :	Distribution System Operator	
PX :	Power Exchange	
NEMO :	Nominated Electricity Market Operator	



	National level (🇧🇪)	European level (🇪🇺)
TSO	Elia	ENTSO-E
NRA	CREG	ACER
DSO	Ores, Sibelga, Fluvius, ...	EU DSO
PX and NEMO	/	In competition: e.g., EPEX spot and Nord Pool

Figure 2.3.: System Operators, Regulators and Power eXchanges in Europe [48].

2.2. System Operators, Regulators and Power Exchanges

In line with the evolution of the European legislation, different entities were established for operating and regulating the electricity networks and markets at both national and European levels, as illustrated in Fig. 2.3.

The Transmission System Operators (TSOs) are entities who carry out the operation, maintenance and development of the electricity transmission network in their given area. Their missions include: i) building and maintaining the electricity network, ii) assuring a non-discriminatory access to the electricity network, iii) guaranteeing the security and quality of electricity supply, iv) ensuring a continuous balance between the total electricity demand and supply, while preventing voltage violations and line congestions, v) promoting efficient and transparent electricity markets. The entity responsible of the Belgian transmission network is Elia, while the entity managing the coordination among the different national TSOs is ENTSO-E, i.e., the European Network of Transmission System Operators. ENTSO-E is involved in the creation of the network codes, and publishes a European ten-year network development plan (TYNDP) every two years, which presents the foreseen developments of the European electricity network in the next 10 to 20 years.

The National Regulatory Authorities (NRAs) are entities politically and financially independent. They are entrusted to guarantee transparency and competitiveness of the electricity markets, while protecting the rights of con-

sumers. In Belgium, the federal entity is called the Commission of Regulation of Electricity and Gas (CREG). The corresponding entity at European level is the Agency for the Cooperation of Energy Regulators (ACER), which ensures that the national regulatory frameworks evolve within the framework of the European energy development policies. ACER publishes an annual market monitoring report, and they are also involved in the creation of the network codes.

The distribution system operators (DSOs), such as Ores, Sibelga or Fluvius in Belgium, are entities who carries out the safe and reliable operation of the medium to low-voltage networks (below 30 kV). They connect end-users, install electricity meters and communicate the metering to the suppliers. The Clean Energy Package establishes an European entity of distribution system operators (EU DSO entity) to increase efficiencies in the electricity distribution networks, and to improve the cooperation of DSOs with ENTSO-E and TSOs.

Power eXchanges (PXs) are anonymous electronic platforms, which trade blocks of energy by matching the offers and bids of multiple actors. PXs complement bilateral contracting, also called over the counter (OTC) markets, where two actors negotiate bilaterally without others knowing the details of the transaction. OTC markets are still larger in volume since market participants can rely on tailor-made contracts and products, but PXs receive an increasing attention as higher electricity wholesale volume is observed [49]. The PXs are known as more transparent as prices and volumes are published through their platforms, while the details of OTC trades remain between the negotiating parties. Practically, the PXs receive orders from various market participants, and have the responsibility for matching and allocating orders in accordance with the relevant participant agreements and regulations. Electricity trading in PXs can start up to a few years before the delivery via the futures market¹ until an horizon shorter than a day via the day-ahead and intraday markets. In Europe, the two largest PXs are: Nord Pool and EPEX SPOT, which have gradually consolidated their dominant positions through different mergers and acquisitions. Note that these two PXs are also designated as Nominated Electricity Market Operators (NEMOs) in Belgium. The NEMO is a certified PX that is allowed to co-organize the coupling of the day-ahead and intraday markets at European level (more details in Section 2.6).

TSOs have the responsibility of balancing the electricity demand and supply.

¹futures market define the long-term standardised products traded in power exchange, while the term forward market is reserved for products that can be tailored-made and traded OTC. Note that the term forward electricity market can cover both types of products [50].

However, the current legislation in Europe prohibits the TSOs from possessing their own generation and consumption means (including storage units), and, thus, are not able to ensure directly by themselves the stability of their network. For this reason, the TSO outsources its balancing responsibility to private entities called Balance Responsible Parties (BRPs), which are presented in the next section.

2.3. Balance Responsible Parties

A Balance Responsible Party (BRP) is an entity appointed at different access points of the electricity grid, i.e., the interface between transmission and distribution levels where injections and/or off-takes are measured, through a contract with the TSO. The contracted BRP has the responsibility to compose a daily balancing schedules of its portfolio on a quarter-hourly basis. Hence, for instance, residential end-users are represented by an electricity supplier, which acts as a BRP for granting the access to the transmission system network. Examples of BRPs in Belgium include suppliers (e.g., Lampiris or Luminus), generation companies (e.g., Electrabel), large industrial users (e.g., Arcelor-Mittal Energy), or even purely financial entities (e.g., Gazprom Marketing & Trading).

Each BRP is thus responsible of maintaining a balance on a quarter-hourly basis within its portfolio. The portfolio of a BRP can be composed of its own generation, its own consumption, and energy exchanges with other BRPs. These exchanges can occur via the electricity markets, from years ahead up to close to real-time. The long-term markets allow BRPs to hedge short-term price risks and uncertainties, while the short-term markets (up to one day before the delivery) allow BRPs to react against changing operating conditions such as unplanned outages, weather-dependent renewable production, the demand level, or network outages. In short-term electricity markets, the exchanges between BRPs mainly take place in the day-ahead and intraday markets, which are managed by Power eXchanges. Due to the ongoing integration of weather-dependent renewable energy sources, BRPs exchange more and more volumes in these markets as the production profiles of weather-dependent renewable energy sources are less uncertain closer-to-real-time. Fig. 2.11 depicts the interactions of the BRPs up to one day before the delivery of electricity.

The Day-Ahead (DA) market is the main short-term arena for trading energy between BRPs. The market is cleared once at 12H00 in the day before delivery (day D-1) based on an auction mechanism, where hourly blocks of energy are traded for the following 12-36 hours. Practically, the day-ahead market

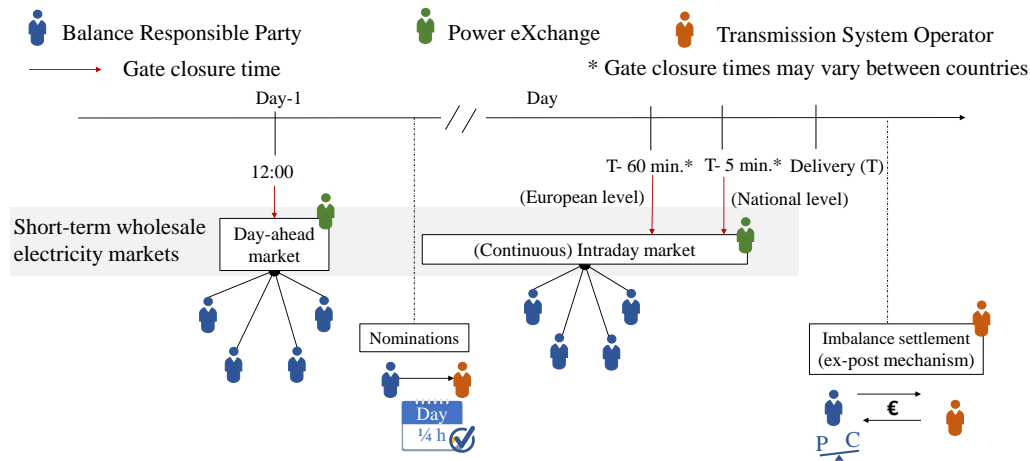


Figure 2.4.: Short-term interactions of the Balancing Responsible Parties.



outcomes are hourly day-ahead prices, which are defined for each hour of the following day (day D), and the volumes exchanged between BRPs. After the day-ahead market clearing, the local TSO requires the BRPs to provide their nominations, i.e., their balanced schedules at both consumption and production levels, on a quarter-hourly basis for the next day. Based on these nominations, the TSO can verify whether the dispatched power units do not lead to voltage violations or line congestions within their control area. Indeed, the day-ahead market-clearing procedure does not fully consider the physics of the transmission and distribution networks, which can lead to impractical market outcomes (more information in Section 2.5). In this case, remedial measures are necessary, which can be i) a change in the grid topology, ii) re-dispatching actions, where the TSO updates the generation and/or consumption scheduling of a BRP, or iii) countertrading actions, where cross-border exchanges are altered by the TSOs. These remedial measures have a cost for the system, which amounts 3.6 billion € for a volume of 61 972 GWh in 2020 [49]. In addition, the TSO's interventions are expected to be more frequent in the near future, due to the improper network representation in the day-ahead market-clearing procedure and the increasing share of weather-dependent renewable energy sources.

After the nominations of BRPs, the intraday market is opened up and allows them to exchange their energy surplus or deficit up to 5 minutes before real-time, which result from forecast errors or unexpected events (as a power plant outage). In Europe, two designs for the intraday market exist: discrete auctions and continuous trading. Both designs can coexist, but continuous trading is the dominant model in Europe as volumes are considered too low for organizing

auctions. In auction-based market, similarly to the day-ahead market, several intraday sessions are organized over the day according to different gate closure times. After each gate closure, the orders are frozen, and are cleared for obtaining the market prices and volumes for the corresponding market period. In a continuous market, matches between BRPs can be performed at any time during the trading session if their respective orders are concordant. Hence, a continuous intraday product is not defined by an unique price common to all BRPs (as in the auction mechanism), but by different prices depending on the time at which the match is concluded. The European Network Codes, and more specifically the Capacity Allocation and Congestion Management guidelines (CACM GL) [51], suggest that continuous trading (wherein an order book aggregating all demand bids and sell offers is continuously updated) should be the main European intraday market mechanism. This offers more trading possibilities for BRPs to adjust their schedules closer-to-real-time, which is important for, e.g., renewable energy technologies, that could be otherwise exposed to imbalance costs.

Indeed, BRPs are charged with imbalance fees each time they deviate from their nominations. These fees are set on a quarter-hourly basis by the imbalance settlement. The imbalance settlement mechanism aims at ensuring that BRPs efficiently support the balance of the electricity system. This mechanism is organized ex-post (after the delivery of electricity) by the TSO, and reflects the cost related to the real-time activation of additional energy (called balancing energy) on BRPs that were not able to honour their schedules. In real-time, a BRP can be in surplus of energy (also denoted as a positive or long imbalance position) or in shortage of energy (i.e., a negative or short imbalance position) compared to its nominations. The net total sum of all BRP's imbalance positions gives the system imbalance, which is managed to be as close as possible to zero. Note that the direction of the imbalance position of one BRP can be different from the direction of the whole system's imbalance position. As the pricing mechanism of the imbalance settlement is not (yet) uniform in Europe, different design choices currently exist [52].

A main design choice is whether the imbalance price is uniquely defined, i.e., "single pricing", or if positive and negative imbalances are priced differently, i.e., "dual pricing". Overall, in both pricing mechanisms, unbalanced BRPs that are in the same direction than the system imbalance pay for the activated balancing energy (red boxes in Fig. 2.5). These two imbalance pricing mechanisms differ when the unbalanced BRPs are in the opposite direction of the system imbalance. For this situation, dual pricing caps the imbalance price at the day-ahead price, while single pricing does not. This difference results

 Balance Responsible Party
  Transmission System Operator










		Dual Pricing System imbalance 		Single Pricing System imbalance 	
		(Energy excess) 	(Energy scarcity) 	(Energy excess) 	(Energy scarcity) 
BRP's Imbalance 	Positive imbalance (e.g., oversupply) 	Highest price of downward balancing energy (MDP) (TSO pays BRP)	Day-ahead price (TSO pays BRP)	Highest price of downward balancing energy (MDP) (TSO pays BRP)	Highest price of upward balancing energy (MIP) (TSO pays BRP)
	Negative imbalance (e.g., undersupply) 	Day-ahead price (BRP pays TSO)	Highest price of upward balancing energy (MIP) (BRP pays TSO)	Highest price of downward balancing energy (MDP) (BRP pays TSO)	Highest price of upward balancing energy (MIP) (BRP pays TSO)

Figure 2.5.: Dual and single imbalance pricing schemes using marginal pricing.

into two types of behaviors for BRPs. In dual pricing, BRPs are incentivized to strictly stick to their energy schedules. Indeed, as shown in Fig. 2.5, dual pricing penalizes, or at least does not reward, any individual imbalances. In single pricing, unbalanced BRPs in the same direction than the system imbalance are charged with an expensive penalty, while the unbalanced BRPs in the opposite direction receive an attractive imbalance price. Indeed, the latter cases of imbalances (green boxes in Fig. 2.5), which provide a real-time balancing service to the system, are remunerated at the same price than the balancing energy activated by the TSO. The European guidelines on balancing markets clearly state that imbalance prices should reflect the real-time value of energy, and favours single pricing over dual pricing (although, dual pricing can be implemented as an exception) [53]. The European guidelines also favour marginal pricing, where the last activated balancing energy bids determine the imbalance price. Overall, the single imbalance price depends i) on the sign of the system imbalance, and ii) on the Marginal Incremental Price (MIP), i.e., the highest price of the energy balancing bids needed for compensating a shortage of energy for a given quarter hour, or the Marginal Decremental Price (MDP), i.e., the lowest price of the energy balancing bids needed for compensating a surplus of energy for a given quarter hour.

In fig. 2.5, it can be observed that the financial responsibility of balancing the system in real time is supported by the BRPs, but it should be recalled that the operational responsibility for maintaining this balance still remains within the hands of the Transmission System Operator (TSO). To achieve this, the TSO buys the needed balancing energy services from other market actors, named Balancing Services Providers (BSPs), through the balancing markets, which is the focus of the next section.




Figure 2.6.: Examples of technical requirements that a balancing service must meet [53].

2.4. Balancing Service Providers

In Europe, the balancing actions, which aim at restoring the system balance, are designed for guarding the system frequency deviation within the range of 49.8 Hz and 50.2 Hz [54]. This is achieved by trading the real-time measured system imbalance from qualified electricity market participants, i.e., the Balancing Service Providers (BSPs).

BSPs are flexible actors, such as gas-fired power plants, demand response facilities, or storage operators, whose power units must meet certain technical requirements for offering standardized balancing services to the TSOs (see Fig. 2.6). The balancing services can be either upward (i.e., an increase in generation and/or decrease in consumption), or downward (i.e., a decrease in generation and/or increase in consumption). The upward balancing services are activated for dealing with negative system imbalances (e.g., undersupply or overconsumption situations causing a frequency drop), while downward balancing services are needed for positive system imbalances (e.g., oversupply situations causing a frequency rise). The costs for activating upward balancing services are usually more expensive than the day-ahead and intraday prices, while the costs for activating downward balancing services are commonly lower than the day-ahead and intraday prices.

As illustrated in Fig. 2.7, the balancing services are procured via two market segments: i) the balancing capacity market, which remunerates the availability of the balancing services during the contracted period, and ii) the balancing

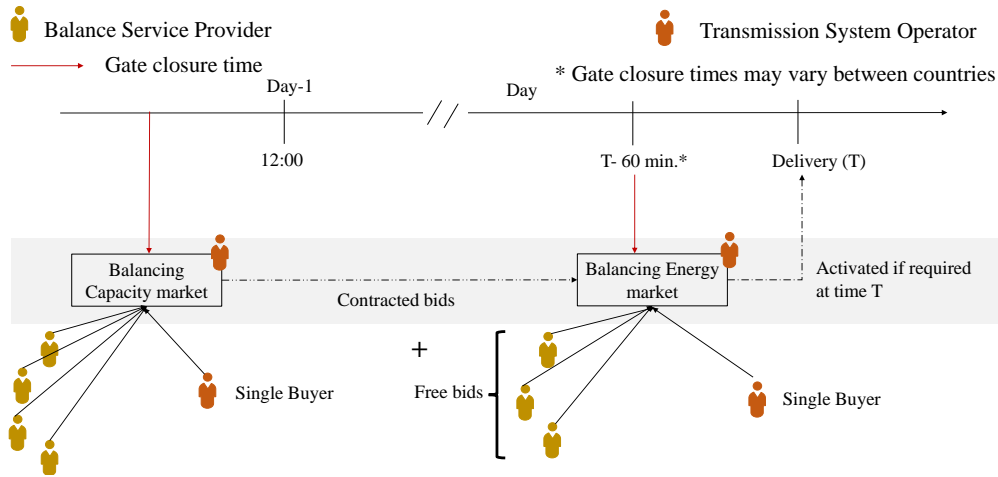


Figure 2.7.: Short-term interactions of the Balancing Services Providers

energy market, which prices the actual activation of the balancing services in real time. Contracted bids in the balancing capacity market are automatically transferred to the balancing energy market, but non-contracted (also called "free") bids can also be added in the balancing energy market. The balancing capacity costs are socialized by TSOs through their transmission network tariffs, but the costs incurred by the real-time activation of balancing reserves are covered by the unbalanced BRPs.

The services in the balancing markets are standardized into four types of products: i) Frequency Containment Reserves (FCR), ii) Automated Frequency Restoration Reserves (aFRR), iii) Manual Frequency Restoration Reserves (mFRR), and iv) Replacement Reserves (RR). Their chronological activation order is shown in Fig. 2.8. As there is only one buyer (the TSO), the total amount of each balancing service is subject to strict European regulations. For instance, in 2019, the volumes of the Belgian balancing services are set to 80 MW for FCR, 145 MW for aFRR and 894 MW for mFRR.

FCR is the first balancing service to be activated. FCR stabilizes the system frequency in a time frame of seconds (usually up to 30 seconds), with an automatic and local activation. For providing FCR, the BSPs must install specific equipments to continuously measure the network frequency, and to quickly adjust their output power for a time period up to 15 minutes. Overall, the FCR volume is dimensioned to handle the loss of the largest generation unit within the control area. Hence, the FCR volume at ENTSO-E level is fixed at 3000 MW, which are distributed between the different control zones according

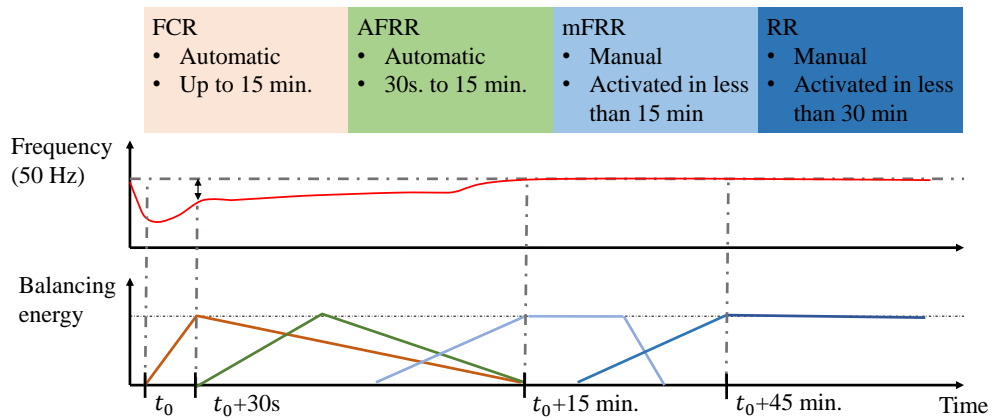


Figure 2.8.: The four types of balancing services in Europe [53].

to their importance in the synchronous area of Continental Europe. After FCR, FRR are the following balancing services to be activated, whose aim is to restore the frequency back to its reference level. More specifically, aFRR denote frequency restoration services that are controlled in an automatic fashion by the TSO, where their activation signal is a set point continuously transmitted by the TSO. The full activation of aFRR is performed in 7.5 minutes, and remains active for the time needed. On the other hand, mFRR necessitate a manual activation from the BSPs, and relieve aFRR in case of a prolonged system imbalance. Finally, RR is the slowest type of services as their activation time ranges from 15 minutes up to hours, and it is used for supporting or releasing FRR activation.

2.5. Clearing the Market

Overall, electricity market prices aim at promoting i) long-run efficiency, i.e., how well efficient investments in the electricity system are encouraged, and ii) short-run efficiency, i.e., how well the available electricity resources are coordinated for maximizing the social welfare [55]. The social welfare can be defined as the sum of the net consumers' surplus and of the producers' profit from trading, which captures a notion of an economic gain considering the entire society.

Most of the European short-term electricity markets (i.e., the day-ahead and balancing markets) are based on an auction mechanism, where the intersection point between the demand and supply curves provides the market price and

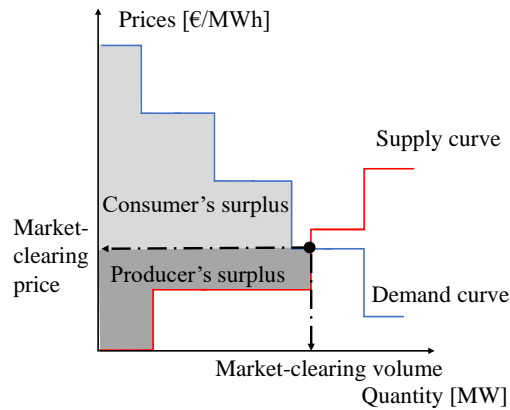


Figure 2.9.: Auction mechanism based on demand-supply curves.

volume. The demand curve is the consumer bids (ranked in a descending order) representing their demand quantities along with their maximum price they are willing to pay. Similarly, the supply curve is the producer offers (ranked in an ascending order), which represents their available production quantities along with their minimum selling prices. An illustration of the auction mechanism based on these demand-supply curves is given in Fig. 2.9.

With this mechanism, all producers with a lower offer price than the market-clearing price are scheduled and paid by the market-clearing price, and, thus, benefit from the difference between the market-clearing price and their operational cost (which enables them to recover their investment costs). Similarly, all consumers with a higher bid price than the market-clearing price pay their energy at the market-clearing price, and, thus, benefit from the difference between the market-clearing price and their utility cost. By finding the equilibrium point between this consumers' and producers' surplus (indicating by the two shaded areas in Fig. 2.9), the market-clearing procedure maximizes the social welfare [56]. In practice, this equilibrium problem can be written as a centralized optimization program, which maximizes the social welfare subject to the producer bids and consumer offers.

Another form of trading applied in the electricity markets is continuous trading. As illustrated in Fig. 2.10, the offers and bids are not simultaneously matched after one gate-closure time, but they are sequentially treated according to the first-come-first-served rule. If the price of a new buy order is equal or higher than the price of an existing sell order (or the opposite), both orders are automatically matched. Each match is settled following the pay-as-bid principle, where the price is set by the oldest order of the two. The unmatched orders are

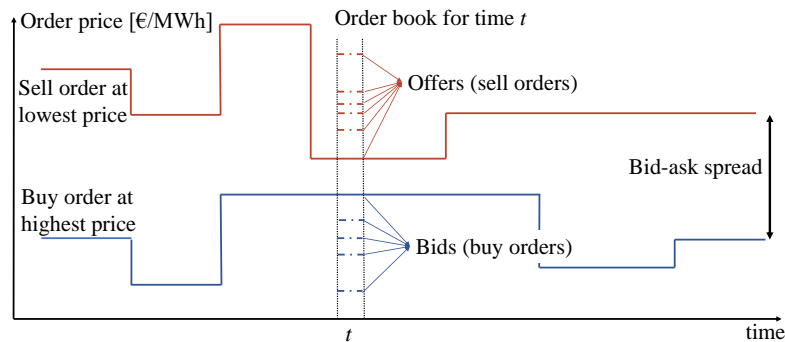


Figure 2.10.: Continuous trading.

stored in an order book, where all their details are accessible to other market participants for facilitating future matching. As depicted in Fig. 2.10, the order book is structured of two sides: the sell and buy orders. The sell orders are sorted in descending order, while the buy orders are sorted in ascending order. The bid-ask spread is the difference between the lowest price of a seller and the highest price of a buyer, which is used as a proxy for measuring the liquidity of this type of market. In a continuous trading scheme, the same quantity of energy for a certain delivery period may have different prices depending on the time of the transaction. The continuous trading scheme is currently the favoured option for the European intraday market.

One of the fundamental difficulties when organizing the electricity markets is that the commercial flows do not necessarily reflect the physical flows. This can be due by i) an inexact representation of the electricity network in the market-clearing mechanism, and ii) an inexact incorporation of the (non-convex) techno-economic constraints of the power units. The less compatible the results of the market is with the physical reality of the system components of the electricity system, the greater will be the need for corrective measures. These corrective measures may entail huge costs for both system operators and market participants.

Concerning the market-clearing mechanisms accounting the electricity network limitations, two major options are currently used worldwide [57]: nodal and zonal pricing, which are showcased in Fig. 2.11. On one hand, nodal pricing, which is typically used in US-style markets, determines local price signals for each grid node. These prices are derived from the balancing equations of each grid node, i.e., the nodal generation must equal nodal consumption, while accounting the import/export of power with other nodes. This necessitates the computations of the physical flows of each line, which is usually achieved by

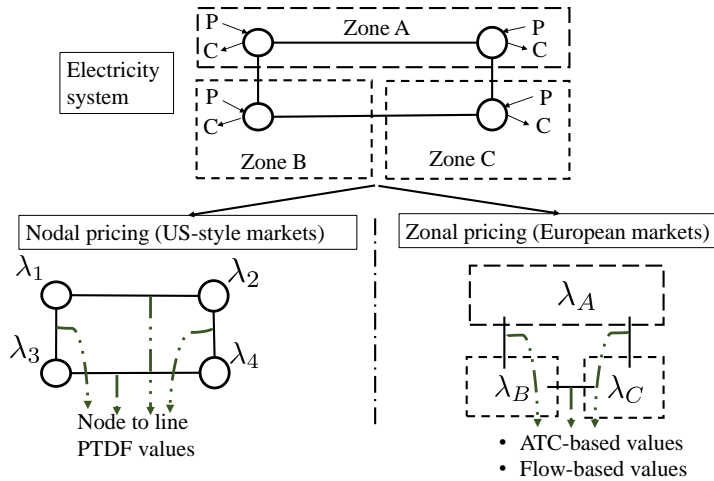


Figure 2.11.: Pricing with network limitations for a 4-node/3-zone network [57].

nodal Power Transfer Distribution Factors (PTDFs). These values provide a linearized approximation of the load flows in response to an incremental change of import/export at each node. Note that other network constraints, such as thermal capacities of the lines, are also incorporated. On the other hand, via zonal pricing, the market-clearing algorithm enforces the balancing equation only at zonal level (which delimits an ensemble of grid nodes), inside which the network is represented as a copper plate (i.e., having an unlimited transmission capacity within the zone). This pricing scheme thus defines an uniform price for the whole zone, but at the expense of a loss of modeling accuracy, which may result in internal congestions. This pricing scheme is currently applied in European electricity markets. The European zones are commonly delimited according to national borders (which thus enabling consistent prices within a same country, such as Belgium), but countries can be subdivided into several zones in order to cope with internal structural congestions (such as Italy). The financial exchanges operated between zones, also referred to as cross-border capacities, can be allocated following two methods: Available Transfer Capacity (ATC) or Flow-Based (FB) methods. The former method limits the maximum commercial exchanges between zones by predefined lower and upper bounds (the ATC values). These bounds are calculated ex-ante (in day-2) by the TSOs, and define the maximum commercial exchanges between zones that are compatible with the physics of the network and operational security standards. However, in the market-clearing procedure, these ATC values are hard constraints, and, consequently, provide conservative cross-border capacity constraints. On the other hand, the latter option, the FB method, introduces zonal PTDFs for characterizing the physical flows between the zones during the market-clearing




Figure 2.12.: Illustration of the linked order (left) and loop order (right).

procedure. These values, which are also determined ex-ante by the TSOs, allow to model linear dependencies between the flows and import/export of each zone during the market-clearing procedure. Evidently, although sharing similarities with nodal PTDFs, zonal PTDFs are less accurate as i) the nodes are grouped into zones, and ii) their values are computed ex-ante based on a reference case (i.e., a forecast of the state of the electricity system at moment of delivery) [57]. Nevertheless, the literature shows that the FB method demonstrates a global increase of trade volumes, which leads to welfare gains compared to the ATC method [58].

While the physics of the network is important, the techno-economic constraints of the power units may also create a dissonance between the market outcomes and the physical reality. In the traditional representation of auction-based markets, the structure of the orders are conventionally defined as price-quantity pairs, where buyers and sellers offer their marginal cost curves. However, such orders cannot reflect all the non-convexities of generation profiles and costs pertaining to the currently existing power units. For instance, thermal power plants have minimum up and down temporal constraints, while their cost structure include starting, operational and shutdown costs. In this setting, trading exclusively hourly price-quantity pairs in the day-ahead electricity market cannot provide the possibility for BRPs to inform the market-clearing mechanism about all their relevant techno-economic constraints, which may induce imbalance penalties in case of infeasible schedules. In light of this shortcoming, block and complex orders have been introduced in the day-ahead market. Fig. 2.12 illustrates two block orders: the linked block order and loop order. The first type of order is defined by parent and child orders, where the acceptance of the child block depends on the acceptance of the parent block. This order allows, e.g., to detail more precisely in the market the starting, shutdown and operating costs of the thermal units. On the other hand, the loop order is a family of two orders that are executed or rejected together. This order was designed for better informing the market about the operation of energy storage systems. However, the use of such orders rises important



Figure 2.13.: The bidding zones in Europe [60].

algorithmic and economic challenges in the European day-ahead electricity markets. Indeed, these orders introduce non-convexities as they are characterized by binary values, which increases greatly the complexity of defining the market equilibrium price. For more detailed discussions on this problematic, the interested reader is referred to the work [59] and references therein.

2.6. Towards an Integrated European Electricity Market

The European approach for merging and coordinating each national electricity market is zonal market coupling. This approach defines bidding zones (usually defined by national borders), where the prices are uniform at each zone and market period. Fig. 2.13 shows the current delineation of the bidding zones in central, west and north Europe.

Integrating the European electricity markets represents both economical and technical advantages. As an order of magnitude, the total economic benefit for integrating the electricity market was estimated at 4 billions € for the year

2016 [61]. In addition, the more an electricity system is interconnected, the more its level of inertia is increased, as more synchronously rotating machines are coupled. This allows mitigating drop or spike of frequency when a mismatch between production and consumption occurs, which facilitates the operation of the network. Another advantage of exchanging electricity across borders is the more effective use of resources that are geographically dispersed. As the wind and solar conditions are uneven across a wide geographical area, the collective behavior of renewable energy source's production profiles can mitigate their intermittency at some extent [62]. Finally, the relative importance of the most severe incident decreases as system size increases, while the balancing capacity energy necessary to cope with such incidents can be mutualized across the different Member States. However, it should be recalled that the resulting price convergence across the Member States creates winning situations (consumers where prices were high can pay less after integration, and producers where prices were low have an access to a market with higher prices) and losing situations (producers with high prices face increased competition, and consumers with low prices experience higher prices). Besides, the needed cooperation for integrating the electricity markets could reinforce the market power of certain actors, which may limit the promised benefits. For instance, in 2014, the two leading European power exchanges, EPEX SPOT and Nord Pool, have agreed to not compete for their short-term electricity trading services across Europe, which has led the European Commission to fine them about 6 million € in cartel settlement [63].

Integrating the Day-Ahead Market

All the bids and offers in European day-ahead markets are coupled via a common optimization algorithm called the Pan-European Hybrid Electricity Market Integration Algorithm (EUPHEMIA), which aims at maximizing the social welfare, while accounting the interconnections between the zones. The operation of the algorithm is performed by the Nominated Electricity Market Operators (NEMOs) on a rotational basis. The EUPHEMIA market-clearing task is arduous as it deals with a wide variety of non-convex orders (e.g., the linked block or loop orders), which forces the implementation of heuristics and iterative solving methods to deal with, e.g., the existence of multiple equilibrium solutions (see reference [64] for a public description of the EUPHEMIA algorithm).

Integrating the Intraday Market

The closer to real time the more specific the needs of the Balance Responsible Parties (BRPs) are and, consequently, the more difficult it is to find the right

counterparty. In a context where intraday volumes are naturally low, continuous trading became the dominant model in Europe for intraday markets. Following the Capacity Allocation and Congestion Management guidelines [51], the Single Intraday Coupling (SIDC) project has delivered in 2018 the first go-live platform of an European continuous intraday market. Through this platform, the BRPs can match their intraday orders with orders submitted outside their zones, as long as enough transmission capacity is available.

Integrating the Balancing Market

The target model for integrating the European balancing markets is a TSO-TSO model, where the cross-border exchanges of balancing services are exclusively handled by TSOs via a common merit order list. The common merit order list ensures that the cheapest balancing energy bids are firstly activated across different zones. Naturally, this integration process requires i) a high level of harmonization of balancing products, and ii) a highly sophisticated IT infrastructure due to the need of robust and efficient real-time communication. The technical requirements of the standard European balancing products are designed to be as much as possible technology neutral (which allows ensuring a level playing field for all technologies), rather than oriented for particular technologies. The different pilots for harmonizing the European balancing markets are shown in Fig. 2.14.

The FCR Cooperation is a regional project involving TSOs from Austria, Belgium, the Netherlands, France, Germany, Switzerland, Slovenia and West Denmark for procuring their FCR in a common market. This FCR Cooperation is performed with daily auctions considering 4 hours-based symmetric products. The auction takes place every day at the gate closure time 08:00 CET, and applies for the next delivery day. In parallel, another cooperative project is the International Grid Control Cooperation (IGCC), which aims at avoiding simultaneous activations of opposite aFRR in adjacent zones. Practically, this imbalance netting process is carried out continuously with a 5-seconds refreshing rate based the imbalance signals of the different coupled market zones, and is limited by the available transmission capacity. Hence, instead of simultaneously activating aFRR in two adjacent zones, only the TSO with the higher absolute system imbalance activates the required difference of aFRR. Note that, between 2012 and 2019, the accumulated benefit from this imbalance netting was estimated at 475 million €, while the investment and operational costs are respectively 20 million € and 1 million € [45].

The PICASSO project (launched in 2017) is focused on establishing the European aFRR platform, which is legally enforced to be operational in July 2022.

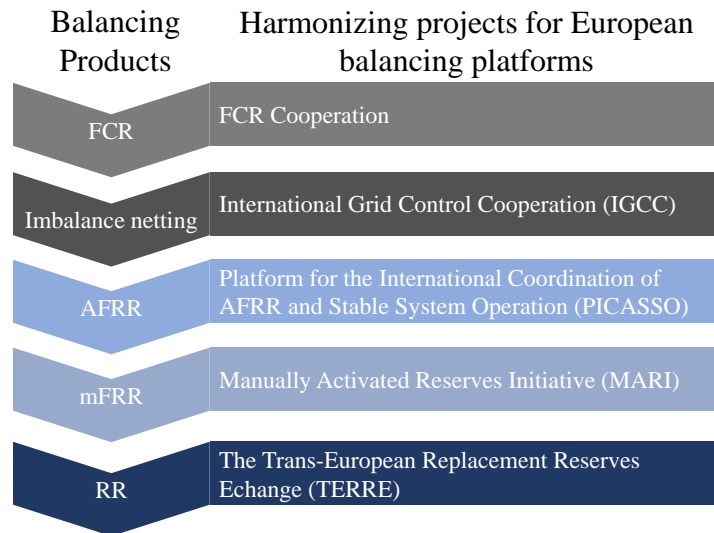


Figure 2.14.: The different balancing pilots for establishing the European Balancing markets.

This platform will allow to exchange standard aFRR balancing products across Europe. In the same vein, the MARI and TERRE projects are respectively responsible for implementing the mFRR and RR European platforms.

2.7. Trading in European Short-Term Electricity Markets

Recalling Section 1.2 of Chapter 1, the main flexibility options for a market actor (either a Balance Responsible Party (BRP), a Balancing Service Provider (BSP), or both), are thermal power plants, demand response, energy storage systems, active control of weather-dependent renewable energy sources, and (cross-border) energy exchanges. Elaborating efficient trading strategies for leveraging their economic value in short-term (European) electricity markets has been an extensive subject of research in the power system community over these two past decades. This Section exposes some key challenges and trading approaches for each market segment, while highlighting the importance of an accurate representation of the real-time market operation in the trading strategies of BRPs and BSPs.

Trading in Day-Ahead Electricity Markets

The day-ahead electricity market was the first short-term market segment receiving a high attention from the power system community. As thermal power plants were formerly the predominant power technologies, early references were focused on developing day-ahead trading strategies that adequately incorporate their (non-convex) operational constraints (e.g., starting up, shutting down, ramping up and ramping down), while assuming a perfect information about the day-ahead electricity prices – see, for instance, the mixed-integer linear programming problem in [65]. The latter hypothesis concerning perfect foresight was gradually lifted in the power systems community, as shown in [66], where crafted scenarios of day-ahead electricity prices are introduced in the trading model of a thermal power plant for better accounting the day-ahead price uncertainty. This notion of uncertainty was even further intensified with the introduction of weather-dependent renewable energy sources since their production profiles are highly intermittent and variable. Such considerations have boosted the energy forecasting community in developing probabilistic forecasting models, which allow better informing the trading strategies of BRPs about forecast uncertainty – see, e.g., [67] and references therein. The interest of such probabilistic forecasting models is illustrated in [68], which derives optimal trading strategies for wind producers based on probabilistic predicted quantiles of their production profiles. In this line, along with the development of probabilistic forecasting models, two-stage stochastic decision-support tools were progressively introduced in the power system community, whose resulting objective function, on the one hand, maximizes the expected profit in the day-ahead electricity market, and, on the other hand, minimizes the expected imbalance penalties that may occur in case of real-time energy deviations [69]. Indeed, as the weather-dependent renewable energy sources were increasingly becoming technologically mature, their granted advantages (e.g., a fixed feed-in tariff or no balancing responsibility) in the electricity markets were loosened, which exposes their owner to imbalance penalties in case of a real-time energy deviation.

Alternatively, with the advent of carbon-free energy resources, the market participation of energy-constrained technologies (such as energy storage systems and demand response) have also received an increasing attention from the power system community. Practically, energy storage systems are additionally constrained w.r.t. thermal power plants by the physical capacities of their reservoir (that can be, e.g., water tanks or battery cells), which reduces their energy availability over time [70]. On the other hand, the energy availability of demand response typically depends on user-specified constraints, which can be: i) a buffer capacity of an industrial process, which allows stopping the activity

for a certain period of time [71], ii) the thermal comfort of end-users, which dictates the behavior of thermostatically loads [72], and iii) the travel routines of drivers, which impose that batteries of electric vehicles are sufficiently charged at certain periods of time [73]. Overall, demand response providers are essentially small, decentralized and self-interested, and their interactions with the day-ahead electricity market are commonly mediated by an aggregator (e.g., a BRP grouping an ensemble of demand response providers). Note that the development of adequate business models defining the interaction between the aggregator and demand response providers is still under research [18], [19]. Overall, when participating in the day-ahead electricity markets, these technologies attempt to capture price spreads by charging during off-peak market periods (characterized by low prices) and discharging during peak market periods (subject to high prices), which allows mitigating temporal mismatches between electricity production and consumption over the day. However, due to the limited controllability of the aggregator on the behaviors of demand response providers, the aggregator has to deal with unexpected energy deviations in real time, which also exposes him to imbalance penalties.

Hence, the real-time market operating conditions, and, consequently, their modeling assumptions, have a direct impact on the day-ahead trading strategies of BRPs. This is illustrated in [74], where a wind power producer adopts more efficient trading strategies when detailing more precisely the real-time market operating conditions in its day-ahead trading strategy. Generally, the modeling of imbalance prices in day-ahead trading strategies can be of gradual complexity: i) a deterministic representation of the imbalance penalties [75], ii) a stochastic representation of the imbalance penalties (e.g., via a set of scenarios) [76], and iii) a representation of the imbalance price formation within the trading strategy of the BRP [74]. In i) and ii), the trading strategies of BRPs are based on the assumption that their imbalance positions would not influence the formation of the imbalance price (i.e., a price-taker approach). On the other hand, the formulation in iii) allows to capture the impact the imbalance position of the BRP on the imbalance price formation (i.e., a price-maker approach). As the imbalance settlement is small in volume by design, the latter approach is selected for our market application as it naturally takes into account the market small size inherent to the imbalance price settlement (see Chapter 4).

Finally, considering the European day-ahead market framework, a large share of developed trading strategies considers only one price signal per market period (which thus implicitly considers a single zone, wherein the electricity network is represented as a copper plate), for which their trading intention is uniquely

provided to the market through price-quantity orders (which are insufficient for informing the market-clearing algorithm of all the techno-economical constraints of the BRP portfolio). However, referring to Section 2.5, failing to properly account the network limitations and the techno-economical constraints of the portfolio may lead to energy schedules infeasible in practice, which may entail costly imbalance penalties for BRPs. Yet, to the best of the author's knowledge, there is still little progress in the power system community for incorporating such considerations in the BRP trading strategies. This may be explained by the non-availability of the European day-ahead market algorithm EUPHEMIA, where only some general aspects are given in the public documentation [64]. The lack of its full mathematical procedure renders the development of a (computationally-efficient) unified European day-ahead electricity market framework a highly difficult task for researchers (e.g., due to the existence of multiple equilibrium solutions) [77]. This mathematical complexity is even further exacerbated when such non-convex market-clearing mechanisms are solved within the trading strategies of BRPs (see, e.g., the model in [78]). This lack of an official European market-clearing model also impacts the studies on the added value of different block orders (e.g., the linked block order or the loop order) over the conventional price-quantity pairs in the trading strategies of BRPs. Indeed, the ex-post performance of block orders w.r.t. conventional ones is currently investigated using either i) already outputted day-ahead price signals by EUPHEMIA [79], with a price-taker assumption (i.e., the trading strategy of the BRP does not influence the market-clearing mechanism), or ii) an ex-post centralized economic dispatch problem [80], which assumes a full disclosure of all the techno-economic constraints of the BRPs to the market operator.

Trading in Intraday Markets

The day-ahead energy schedules of Balance Responsible Parties can be refined over the day using the intraday market opportunities. Such opportunities are commonly studied considering a Virtual Power Plant (VPP), which is a market actor combining different technologies in its portfolio (e.g., a wind power unit complemented with an energy storage system) [29], [81]–[83]. This extension towards the intraday stages for VPPs results from i) the decrease over time of the uncertainty pertaining to the production profiles of weather-dependent renewable energy sources, and ii) the necessary continual monitoring of the energy availability of energy-constrained units (e.g., energy storage systems or demand response), whose scheduling decisions can be adapted over the day. Overall, the most challenging part lies in the multiple sequential intraday stages, where information can be updated at each stage, which increases drastically the dimensionality, and, thus, the tractability of the associated decision-support tool. Different trading strategies have been developed in

accordance with both auction-based and continuous intraday formats.

Considering the auction-based format (which is, e.g., well-anchored in the Iberian electricity market [69]), rolling optimization (also denoted as model predictive control or receding control horizon in the literature) approaches are generally used for readjusting the energy schedules of VPPs [29], [81]. In this setting, the dynamic updates of forecasts are leveraged by solving an intraday decision-support tool (with a look-ahead horizon) recursively over the day. More specifically, the intraday decision-support tool is formulated and solved at each decision stage with a look-ahead horizon, but only the trading decisions in the first time step of the look-ahead horizon are implemented. Then, for the next stage, the look-ahead horizon is shifted forward, and the calculation is repeated, based on updated forecasts and previous trading decisions. At the end of the day, the final decisions of the market actor are thus the sequence of the decisions in the first time step of each look-ahead horizon [84]. This framework is then generally extended to a stochastic process, where the uncertainty of future electricity quantities (e.g., the intraday prices or wind power) is mitigated via scenarios.

On the other hand, the continuous intraday format increases even further the dimensionality complexity of the intraday trading strategy. Indeed, in this format, the market actor can continuously submit its bids and/or offers, while observing the trading opportunities contained in the order book. Hence, the buy and sell prices of a delivery may evolve over time, while, in the auction-based format, the intraday market is cleared at different gate-closure-times over the day. Due to these particularities, reinforcement learning approaches have been preferred over stochastic optimization procedures in [85], [86] for participating in the continuous intraday market. Note that, nevertheless, a multistage stochastic optimization approach is still used in [87], where advanced decomposition techniques and dedicated sampling approaches have been leveraged for countering the curse of dimensionality.

Trading in Balancing Markets

The balancing market opportunities are typically studied in the day-ahead stage by the power systems community, along with the opportunities offered by the day-ahead electricity market. This renders the market participant as, on one hand, a Balancing Service Provider (BSP) and, on the other hand, a BRP. For simplicity reasons, the trading strategy in both day-ahead electricity and balancing markets are commonly jointly optimized [88]–[90], even though their auctions occur separately in Europe. Such joint optimization scheme aims at identifying the best trade-off between the revenue streams coming from i) the

day-ahead market, ii) the balancing capacity market, and iii) the balancing energy market.

One of the most challenging part when proposing balancing services in day-ahead is to properly consider the uncertain activation of balancing energy in real time. This issue is particularly exacerbated for energy-constrained technologies (e.g., energy storage system or demand response), which is not only constrained by the installed capacity (MW) and ramping abilities (MW/min) (as is the case of thermal power generation), but also by their energy availability (MWh). Failing to properly anticipate the impact of the uncertain activation of scheduled balancing services on their energy content may thus result into overly optimistic bidding strategies, i.e., offering balancing capacity that the market actor is unable to deploy when requested in real time. This may result in costly real-time balancing penalties and/or a temporary exclusion from participation in balancing markets. The representation of the real-time activation of balancing energy can become rapidly computationally-intensive for the optimization problem. For example, if one considers 3 activation values for each of the 96 time steps of the day, one will end up with 3^{96} possible trajectories to feed its stochastic decision-support tool. The balancing activation is thus considered in [91] using an average activation rate at each time step. However, this deterministic procedure does not guarantee that the resulting scheduling is feasible when the actual balancing activation deviates from its mean value. In the same vein, the balancing energy requirement is satisfied in expectation in [92], where the anticipated activation of balancing energy is represented through different scenarios. For alleviating this issue, the work in [93] ensures the real-time feasibility of balancing activation by considering a worst-case scenario in the day-ahead stage, whereas reference [90] considers probabilistic constraints to account for uncertain balancing activation.

Overall, as the rising shares of weather-dependent renewable energy resources introduce a higher degree of uncertainty in real-time market operating conditions, the role of real-time markets (i.e., the balancing energy market and imbalance settlement) become central in forming efficient trading strategies [94]. In line with this evolution, the European balancing market favours the single price imbalance settlement, which better reflects the real-time value of energy. This imbalance settlement mechanism allows BRPs to provide real-time balancing services, which are presented in the next Section.

2.8. Provision of Real-Time Balancing Services

The European balancing products are activated based on real-time frequency deviations measurements. This reactive design implicitly assumes that no anticipated future system imbalances are alleviated by the TSOs. This design generates high spikes in the balancing energy prices during periods where high system imbalances are measured. Recalling Section 2.3, the resulting balancing activation costs are then transferred to unbalanced BRPs via the single price imbalance settlement.

For each imbalance settlement period (typically 15 minutes), each BRP is thus defined by an imbalance position, which is priced by the single imbalance price. This imbalance price provides incentives (financial rewards) to Balance Responsible Parties (BRPs) when they contribute to restore the system balance by their imbalance position (e.g., a deficit of energy in the power system and a surplus of generation for the BRP). Recalling Fig. 2.5 in Section 2.3, the single imbalance price is based on i) the direction of the system imbalance, and ii) the cost of the most/least expensive activated upward/downward balancing energy bid. Based on this marginal pricing method, Fig. 2.15 depicts how the intentional imbalance position of a BRP can influence the single imbalance price. A negative system imbalance ($SI < 0$), i.e., the case a) in Fig. 2.15, defines a generation shortage in the system, while a positive SI, i.e., the case b) in Fig. 2.15, reflects a generation surplus. In case of generation shortage ($SI < 0$), the TSO activates upward reserves, which drives the imbalance price towards a high price regime (typically exceeding the day-ahead prices). In this case, a negative imbalance position aggravates the shortage of the system. Hence, the BRP in this negative imbalance position would have to pay an expensive price to the TSO for the additional activation of upward reserves. If the BRP has a positive imbalance position, its imbalance decreases the overall balancing needs. In that case, the unbalanced BRP would receive an attractive price for each ‘excess’ MWh. On the other hand, in case b) of Fig. 2.15, the system is in generation surplus ($SI > 0$). Downward reserves are thus activated, thereby lowering the imbalance price (typically leading to an imbalance price below the day-ahead energy market prices). Using the same reasoning, we can observe that the single imbalance pricing in this setting incentivizes financially BRPs to adopt a negative imbalance position.

Generally, the BRPs that degrade the system balance are penalized, while those that help maintaining the system balance are rewarded. Hence, the BRP, by adopting an intentional positive (e.g., by overproducing) or negative (e.g., by overconsuming), can provide real-time balancing services to the system outside the standard submission of balancing energy bids. **These opportunities**

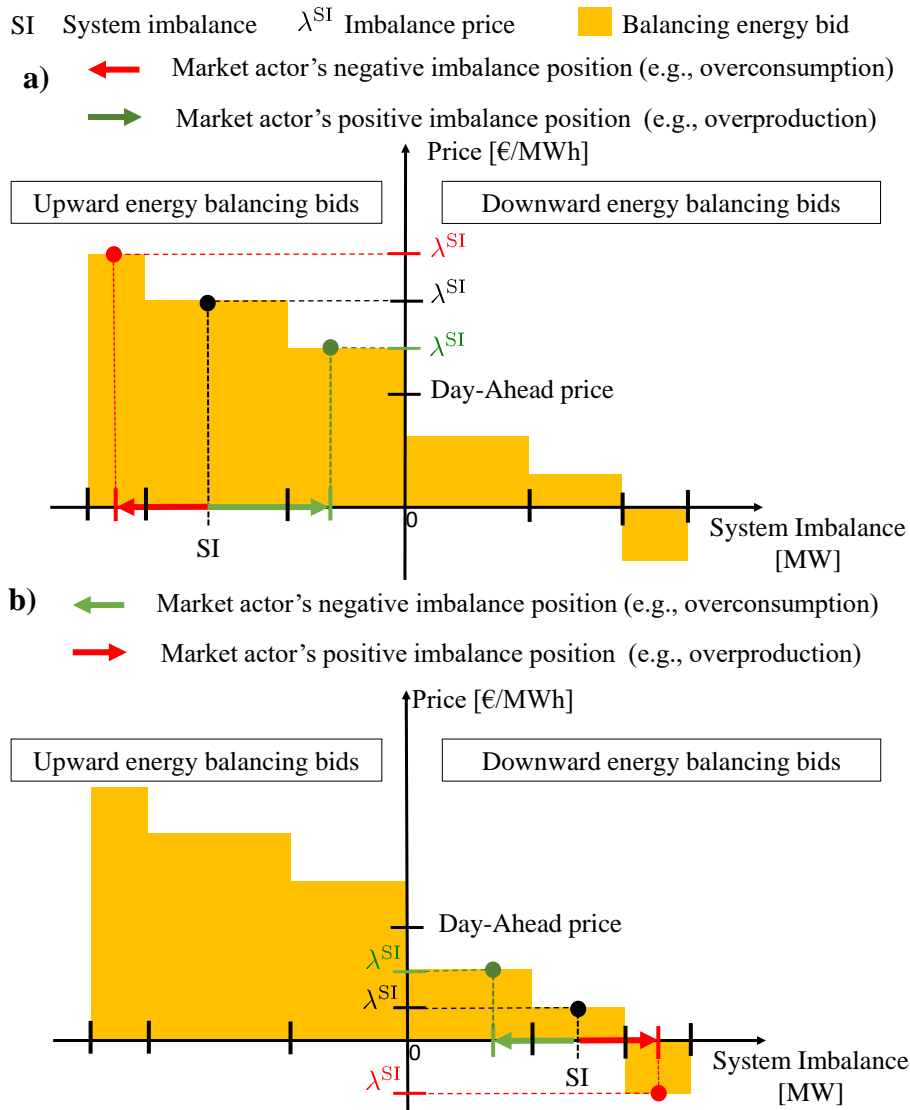


Figure 2.15.: a) Impact of the imbalance position of a market actor in case of a negative system imbalance (e.g., shortage of energy). b) Impact of the imbalance position of a market actor in case of a positive system imbalance (e.g., oversupply of energy).

complement the European balancing markets by fostering the use of flexible resources that do not necessarily meet the technical requirements of standard balancing products. These real-time balancing services can be also referred in the literature by passive balancing (in reaction to the balancing services that are actively activated by the TSO), or even implicit balancing. The provision of such real-time balancing services are the focus of this report, and they are studied based on the Belgian power systems. The Belgian power system is a good proxy for illustrating the aspired integrated European electricity markets as i) the imbalance settlement period is set to be 15 minutes, ii) the Belgian TSO (Elia) employs a paid-as-cleared remuneration for balancing energy, iii) the imbalance prices are set at the marginal activation cost of the balancing energy merit order curves, and iv) a single imbalance pricing method is implemented.

For reaching the full benefit of these real-time balancing services, the publication of timely and comprehensive real time market data is a prerequisite. In this context, e.g., the Belgian TSO publicly provides day-ahead information with near real-time updates on the available balancing energy offers (capacity and activation costs), and the system imbalance position for each 15-minute period [95]. Such information allow BRPs and Balancing Services Providers (BSPs) to better represent future balancing energy market-clearing procedures and associated imbalance price formations in their trading strategy. As we can observe in Fig. 2.15, an important variable of the single imbalance pricing scheme is the total net system imbalance, which directly affects the imbalance price regime (i.e., either low for positive system imbalance and high for negative ones).

Hence, for providing these real-time balancing services, several authors have studied near real-time dispatch strategies based on the system imbalance signal. Abdisalaam et al. [37] propose a real-time re-optimization model for assessing the value of residential demand response, while assuming a perfect foresight of the system imbalance. Zapata Riveros et al. [38] propose a deterministic optimization approach for an aggregator of residential cogeneration systems, where econometric models predict the future system imbalance and imbalance price signals. Koch presents in [39] a trading strategy, based on a logistic regression model, aiming at capturing the arbitrage value between the intraday market and the single price imbalance settlement. In all of these works, the procurement of real-time balancing services relies on deterministic decision-support tools, in which optimal decisions, i.e., the intentional imbalance positions of the actor, are based on a single-point forecast of the system imbalance direction (which is positive or negative). **In this dissertation, we go beyond these deterministic approaches by, in a first step, developing probabilistic**

time series forecasting models of the system imbalance that, in a second time, will feed risk-aware decision-support frameworks tailored for providing real-time balancing services.

2.9. Conclusion

This Chapter aims at providing an overview of the European short-term electricity markets from a market actor's perspective. Firstly, the driving forces behind the organization of the European electricity markets were commented by detailing each European legislative energy packages from 1996 to date. This novel economic organization of the European electricity system triggered the emergence of new entities. In this line, the roles of Transmission System Operators, National Regulation Authorities, Distribution System Operators, and Power eXchanges, which currently operate and regulate the electricity networks and markets, were presented at both national and European levels. Then, focusing on the European short-term electricity markets, the responsibilities and trading opportunities of the two major European market entities, i.e., the Balance Responsible Party and the Balancing Service Provider, were discussed. In the following, the different market-clearing mechanisms of the short-term European electricity markets were detailed, while presenting how the network constraints and non-convexities in costs and generation profiles are handled in the European framework. Next, a status of the harmonization process of the European short-term electricity markets is provided, which is followed by a brief overview of the challenges and trading approaches for each market segment, highlighting the importance of the real-time electricity markets. Finally, the market opportunities provided by the provision of real-time balancing services are explained.

In the next chapter, we focus on developing time series probabilistic forecasting methods for accurately anticipating future system imbalances, which is a key variable for providing appropriate real-time balancing services.

CHAPTER 3.

Towards High-Quality Probabilistic Time Series Forecasting Using Neural Networks

Knowing future trends of key electricity quantities (e.g., electricity demand or prices) is crucial for successful trading strategies in a competitive environment [68]. This is why the development of time series forecasting tools is abundant in the power systems literature since the introduction of electricity markets. Currently, time series forecasting models derived from deep neural networks achieve top prediction performance [96], [97]. In this line, this Chapter focuses on neural architectures of gradual complexity for predicting system imbalances. More specifically, this Chapter investigates how the flexible nature of neural networks can be leveraged for better capturing the complex time dependencies between both past observed and future known information.

This Chapter is organized as follows. A general introduction of time series forecasting is provided in Section 3.1. Section 3.2 details the forecasting application, which concerns the prediction of system imbalances. Next, two architectures of neural networks, i.e., the feed-forward and recurrent neural networks, are presented in Section 3.3 and Section 3.4, respectively. Section 3.5 presents how recurrent neural architectures can be tailored for fostering their accuracy performance in presence of past observed and future known input variables. This leads to the development of sequence-to-sequence recurrent neural networks with and without attention mechanisms. Then, the training and inference stages of neural networks are succinctly exposed in Section 3.6. This is followed by Sections 3.6-3.9, which respectively detail the benchmark methods, the probabilistic error metrics and the case study. Sections 3.6-3.9 allow quantitative observations of the presented neural architectures with respect to other well-established forecasting methods (e.g., Autoregressive moving average (ARMA) model or tree-based ensemble methods), based on real-life market data extracted from the Belgian power systems. Finally, Section 3.10

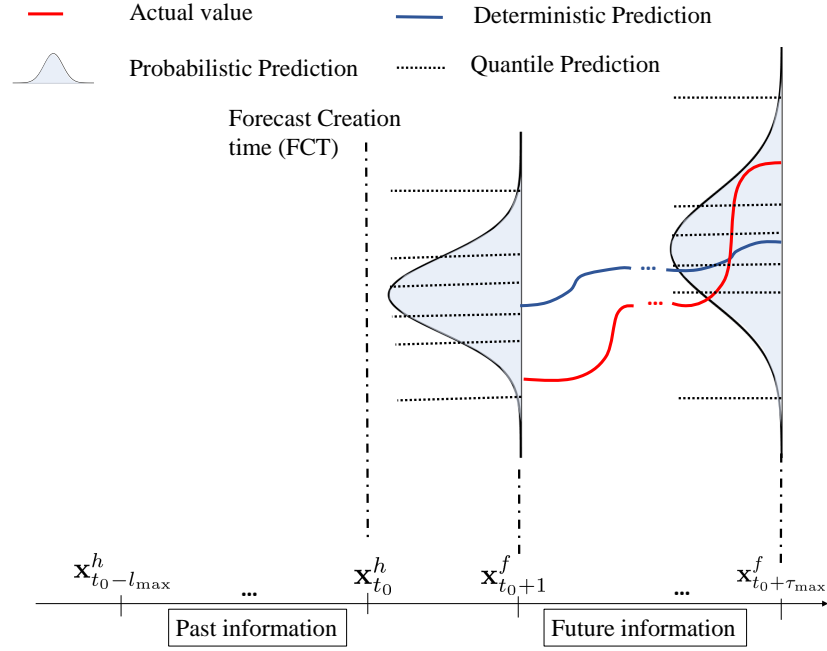


Figure 3.1.: Deterministic and probabilistic time series forecasting

concludes this Chapter.

3.1. Time Series Forecasting

Time series forecasting aims at predicting estimates of a time series at different time points in the future. When applied on electricity time series, this information can then be used in decision-making processes of market actors for supporting their trading strategy. As showcased in Fig. 3.1, a deterministic approach only provides a single time trajectory of the future expected behavior of the time series. The deterministic forecasting approach can be formulated as the following regression problem:

$$\mathbb{E}(y_{t_0+1}, \dots, y_{t_0+\tau_{\max}} | \mathbf{x}_{t_0-l_{\max}}^h, \dots, \mathbf{x}_{t_0}^h, \mathbf{x}_{t_0+1}^f, \dots, \mathbf{x}_{t_0+\tau_{\max}}^f) \quad (3.1)$$

where t_0 is the forecast creation time, and $y_{t_0+1:t_0+\tau_{\max}}$ is the time series to predict over the time steps $\{t_0 + 1, \dots, t_0 + \tau_{\max}\}$ of the prediction horizon. The indices $\{l_{\max}, \tau_{\max}\}$ are respectively determining the number of look-back and look-ahead time periods. The inputs $\mathbf{x}_{t_0-l}^h \in \mathbb{R}^{m_h}$ are past time series observed

(e.g., recent forecast errors of renewable generation) at time step $t_0 - l$, while the inputs $\mathbf{x}_{t_0+\tau}^f \in \mathbb{R}^{m_f}$ are future information, e.g., the prices cleared at the day-ahead stage or calendar information, already known at time step $t_0 + \tau$ of the prediction horizon.

While point predictions can provide some insights on the future, capturing the uncertainty associated with the prediction outcomes can improve the decisions of decision-support tools [98]–[101]. The uncertainty may originate from i) noises in the explanatory variables, and ii) a model misspecification (e.g., a linear forecaster attempting to model a non-linear dependency). The probabilistic forecasting approach is usually written as the following time series regression problem:

$$p(y_{t_0+1}, \dots, y_{t_0+\tau_{\max}} | \mathbf{x}_{t_0-l_{\max}}^h, \dots, \mathbf{x}_{t_0}^h, \mathbf{x}_{t_0+1}^f, \dots, \mathbf{x}_{t_0+\tau_{\max}}^f) \quad (3.2)$$

For predicting such conditional probabilistic distributions $p(\cdot)$ of the variable of interest, two distinct approaches can be found in the literature [102]: i) a two-step procedure, which is based on the addition of a probabilistic forecast error model on top of deterministic forecasts [103]–[106], and ii) methods that directly provide probabilistic predictions of the variable of interest [107]–[109]. In the latter approach, in a non-parametric setting, two methods have been widely applied in the power systems literature, namely kernel density estimators [110], [111] and direct quantile regression models [112], [113]. In this Chapter, the focus is on quantile regression methods, which directly provides the specified q -quantiles $\hat{y}_{t_0+\tau, q}$ such that:

$$q = \mathbb{P}(y_{t_0+\tau} \leq \hat{y}_{t_0+\tau, q}), \quad \forall \tau \in \{1, \dots, \tau_{\max}\}, \forall q \in Q \quad (3.3)$$

where Q is the set of quantiles to predict.

Overall, the deterministic (3.1) and probabilistic (3.2) regression problems are traditionally solved through a supervised learning approach. The supervised approach fits a forecasting function $f_{\Theta}(\mathbf{x}_{t_0-l_{\max}:t_0}^h, \mathbf{x}_{t_0+1:t_0+\tau_{\max}}^f)$ based on a paired dataset containing, on the one hand, the past observed and future known inputs and, on the other hand, the targets. In particular, an error function \mathcal{L} is used for measuring the compatibility of the forecasting function f_{Θ} with respect to all the samples $i \in \mathcal{I}$ of the paired dataset $(\mathcal{Y}_i, \mathcal{X}_i)$, where $\mathcal{Y}_i = (y_{i, t_0+1:t_0+\tau_{\max}})$ and $\mathcal{X}_i = (\mathbf{x}_{i, t_0-l_{\max}:t_0}^h, \mathbf{x}_{i, t_0+1:t_0+\tau_{\max}}^f)$ is the i -sample of the dataset. Based on this error signal, the parameters Θ of the forecasting function f_{Θ} are optimized




Figure 3.2.: Pinball loss for the quantiles $q \in \{0.05, 0.5, 0.95\}$

as follows:

$$\Theta^* = \operatorname{argmin}_{\Theta} \sum_{i \in \mathcal{I}} \mathcal{L}(\mathcal{Y}_i, f_{\Theta}(\mathcal{X}_i)) \quad (3.4)$$

Let T the length of the prediction horizon and $\hat{y}_{t_0+\tau}$ a deterministic prediction outcome, the error function \mathcal{L} is typically the mean square error function in a deterministic framework. For a single sequence, the mean square error function is written as:

$$\mathcal{L} = \frac{1}{T} \sum_{\tau=1}^{\tau_{\max}} (y_{t_0+\tau} - \hat{y}_{t_0+\tau})^2 \quad (3.5)$$

In a quantile regression framework, the error function \mathcal{L} is usually the pinball loss function (also denoted as the quantile loss function), which is expressed as:

$$\mathcal{L} = \frac{1}{T \cdot |Q|} \sum_{\tau=1}^{\tau_{\max}} \sum_{q \in Q} q \cdot \max(0, y_{t_0+\tau} - \hat{y}_{t_0+\tau,q}) + (1-q) \cdot \max(0, \hat{y}_{t_0+\tau,q} - y_{t_0+\tau}) \quad (3.6)$$

where the operator $|\cdot|$ stands for the cardinality of the associated set, and $\hat{y}_{t_0+\tau,q}$ is the q -quantile of the predicted distribution at time $t_0 + \tau$.

Overall, at each time step $t_0 + \tau$ of the prediction horizon, the deterministic forecasting model provides a 1-dimensional output, while the probabilistic forecasting model provides a $|Q|$ -dimensional output describing the uncertainty around the prediction. Fig. 3.2 illustrates the pinball loss function for a specific time step and $q \in \{0.05, 0.5, 0.95\}$. We can observe that the pinball

loss applies asymmetric weighting, i.e., q and $(1-q)$, to the absolute value of the distance between the forecasted q -quantile and the actual observation. Hence, when the forecasted q -quantile is higher than the actual observation, the absolute difference is multiplied by a factor $(1-q)$. On the other hand, when the forecasted q -quantile is lower than the real value, the error is multiplied by a factor q . In this way, in accordance with the quantile definition in Eq. (3.3), the pinball loss penalizes lower (higher) quantiles more severely in case of an overestimation (underestimation).

In general, four types of forecasting models are proposed in the literature [114]: i) persistence methods, which simply consider that future observations will have the same values as the current instance, ii) physical methods, which are based on a detailed mathematical description of the environment governing the variable of interest (e.g., full modeling of the market rules and participant's behaviors to predict market prices), iii) econometric models, which build mathematical models of predefined complexity based on statistic inferences (e.g., autocorrelation), and iv) machine learning methods, which are based on generic, non-linear models that are trained using a self-learning procedure (without being explicitly programmed to achieve the prediction task, and with no arbitrary assumptions on the model complexity). Overall, persistence methods are very naïve, and do not provide useful information for decision-making. Yet, they can be a surprisingly solid benchmark. Physical models, on the other hand, are characterized by a high computational complexity, which hinders their practical utilization for forecasting applications with frequent updates. Moreover, they require detailed analytical formulations for modeling the predicted phenomenon, and are thus not easily transposable on other forecasting applications. Econometric models mostly combine linear inferences of lagged values of the variable of interest, which prevents them from capturing non-linear dependencies present in the dataset. For these reasons, machine learning approaches have recently topped the prediction performance, and this trend is expected to be further intensified by the increase availability of reliable databases and computer capabilities. This Chapter is thus focusing on machine learning forecasting models, with an emphasis on neural models. Different neural architectures of gradual complexity will be presented and tested based on the prediction application of system imbalances. The following Section further details this application and its interest for the electricity system.

3.2. Forecasting System Imbalances

The definition of the term 'system imbalance' may refer to distinct imbalance signals in the literature. Three signals describing an imbalance situation may

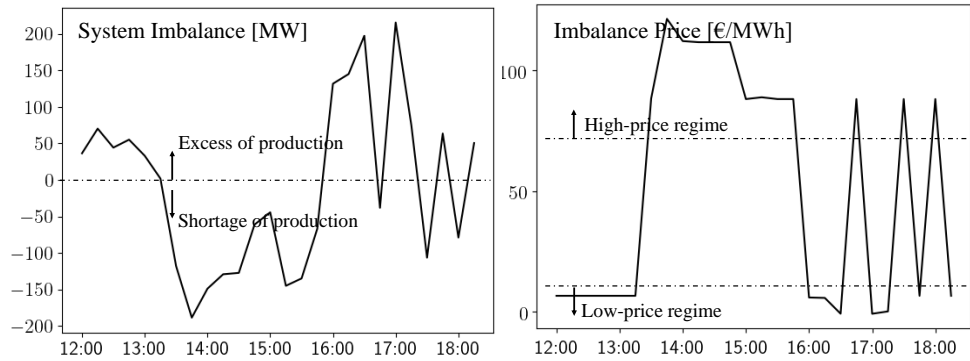


Figure 3.3.: Belgian system imbalances (which is given by the opposite value of the Net Regulation Volume) and imbalance prices on the 1st January 2019.

coexist in Europe: the Area Control Error (ACE), the Net Regulation Volume (NRV), and the residual system imbalance of the control area. The ACE is the difference between the scheduled and actual values of the power exchanged in the TSO control area, i.e., the system imbalance as if no balancing control was performed by the TSO. The net regulation volume (NRV) is the net total balancing volume of energy (upward and downward) activated by the TSO. The residual system imbalance is given by the subtraction between the ACE and NRV signals (which is not equal as the balancing energy markets are reactive by design). In this dissertation, we assume that the TSO has perfectly responded to the actual system imbalance, and, thus, we focus on predicting the NRV signal. Hence, in the following, the term ‘system imbalance’ therefore refers to the opposite value of the NRV signal. This assumption is perfectly aligned with the Belgian imbalance pricing scheme over the years 2016-2019, where the imbalance price regime (either low via the Marginal Decremental Price (MDP) or high via the Marginal Incremental Price (MIP)) was established by the sign of the NRV signal (positive or negative) [115].

Globally, these three imbalance signals depend on the real-time operating conditions of the whole electricity system, which consequently render them highly variable and uncertain. Fig. 3.3 describes such uncertainty and volatility for the Belgian system imbalances on the 1st January 2019. Overall, four factors mainly drive the system imbalance signals, which are: i) the human behaviors, which can lead to a non-rational dispatch of resources in electricity markets, ii) the chaotic dynamics of atmospheric systems, which generate forecasting errors of, e.g., weather-dependent production and electric heating

consumption, iii) unexpected outages of voltage lines or power plants, and iv) an inadequate design of the electricity markets w.r.t. the physical reality of the electricity system (e.g., the hourly time resolution of the day-ahead electricity market yielding ramping trajectories for power plants, which cause imbalances). These four factors are highly challenging to capture, which complicates the prediction of system imbalance signals in comparison with other time series (e.g., wind power or day-ahead electricity price forecasting).

Once predicted, the system imbalance signal can then be used for extracting some information on the imbalance price regime. This dependency is shown in Fig. 3.3, where we can observe that imbalance prices (right figure) are i) high (i.e., prices higher than the day-ahead electricity prices) in periods of energy shortage at system level, and ii) low in periods of excess of production. While the predictions of the imbalance signals can be of value for both Transmission System Operator (TSO) and the Balance Service Providers (BSPs), this report illustrates their usefulness for Balancing Responsible Parties (BRPs), where the probabilistic forecasts of system imbalances will be subsequently used in tailored risk-aware stochastic decision-support tools in Chapters 4 and 5.

3.3. Feed-forward Neural Networks

The feed-forward neural network (FFNN) is the traditional neural network architecture used in the literature [116]. Let $\hat{\mathcal{Y}} = (\hat{y}_{t_0+1:t_0+\tau_{\max}, \forall q \in Q})^\top \in \mathbb{R}^{T \cdot |Q|}$ and $\mathcal{X} = (\mathbf{x}_{t_0-l_{\max}:t_0}^h, \mathbf{x}_{t_0+1:t_0+\tau_{\max}}^f)^\top \in \mathbb{R}^{(l_{\max} \cdot m_h) + (\tau_{\max} \cdot m_f)}$ respectively be the probabilistic predictions of the targeted time series and its associated inputs. In a concise fashion, the forecasting function of the FFNN consists in two linear transformations, with a non-linear activation in between:

$$\hat{\mathcal{Y}} = W_2^{\text{FFNN}} f^{\text{NL}}(W_1^{\text{FFNN}} \mathcal{X} + b_1^{\text{FFNN}}) + b_2^{\text{FFNN}} \quad (3.7)$$

where $W_1^{\text{FFNN}} \in \mathbb{R}^{D \times |\mathcal{X}|}$, $W_2^{\text{FFNN}} \in \mathbb{R}^{T \cdot |Q| \times D}$, $b_1^{\text{FFNN}} \in \mathbb{R}^{D \times 1}$ and $b_2^{\text{FFNN}} \in \mathbb{R}^{T \cdot |Q| \times 1}$ are weight parameters defining the forecasting function, f^{NL} is the element-wise application of a non-linear activation function, and $|\cdot|$ stands for the cardinality of the associated vector. The parameter T defines the length of the prediction horizon. Note that the dimension D , defining the number of processing units in the hidden layer, is a task-dependent hyperparameter to be tuned.

As illustrated in Fig. 3.4, the set of inputs \mathcal{X} forming the input layer is provided to the first hidden layer of the network. The flow of information is then propagated through all the hidden layers to the output layer. More

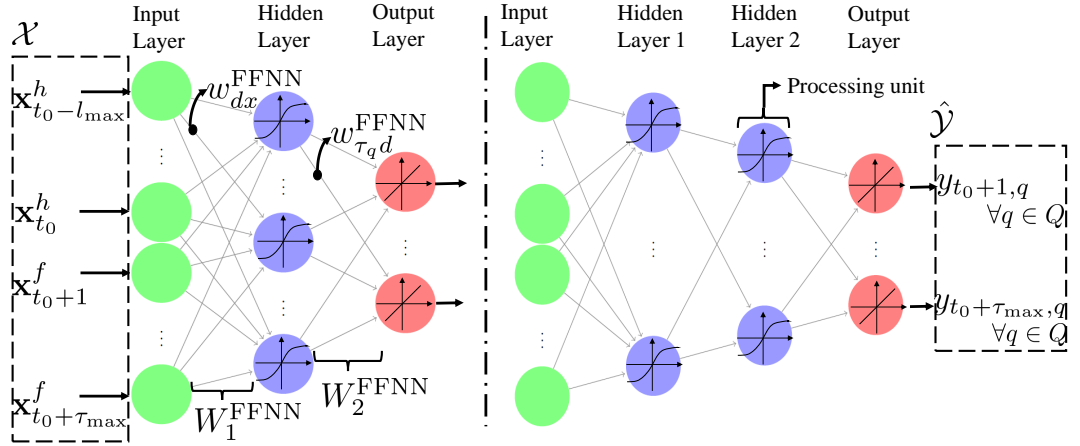


Figure 3.4.: 1-layer and stacked feed-forward neural networks (where biases are omitted).

precisely, the flow of information outputting the prediction $\hat{y}_{t_0+\tau,q}$ is decomposed as follows (biases are omitted):

$$a_d = \sum_{x \in \mathcal{X}} w_{dx}^{\text{FFNN}} x, \quad \forall d \in D \quad (3.8a)$$

$$b_d = f^{\text{NL}}(a_d), \quad \forall d \in D \quad (3.8b)$$

$$\hat{y}_{t_0+\tau,q} = a_{\tau_q} = \sum_{d \in D} w_{\tau_q d}^{\text{FFNN}} b_d, \quad \forall \tau \leq T, \forall q \in Q \quad (3.8c)$$

where $w_{dx}^{\text{FFNN}} \in W_1^{\text{FFNN}}$, $w_{\tau_q d}^{\text{FFNN}} \in W_2^{\text{FFNN}}$ are i) the weight from the input $x \in \mathcal{X}$ to the processing unit $d \in D$, and ii) the weight from the processing unit d to the prediction outcome $\hat{y}_{t_0+\tau,q}$, respectively. The vector $\{a_d, \forall d \in D\}$ results from the weighted sum of the inputs \mathcal{X} by the parameters W_1^{FFNN} , and the vector $\{b_d, \forall d \in D\}$ is given by the non-linear transformation $f^{\text{NL}}(\cdot)$ of $\{a_d, \forall d \in D\}$. Note that, in a regression framework, a linear activation function is often privileged for the output layer, i.e., a_{τ_q} is linearly transformed.

The FFNN architecture is theoretically known to be able to fit any data generating process between inputs and outputs. Relying on the constant increase of computing power, this ability can be augmented by increasing the number of successive non-linear transformations in their architecture, which allows the network to extract more representative features from raw data. Practically, this is usually done by simply stacking several hidden layers in the FFNN architecture, as shown in the right side of Fig. 3.4. In this setting, two hyperparameters are tuned: i) the number of hidden layers, i.e., the number of

non-linear transformations $f^{\text{NL}}(\cdot)$ in Eq. (6.3), and ii) the number of neurons per hidden layer. For both FFNN networks, the forecasting function is fitted by optimizing the weights W^{FFNN} and biases b^{FFNN} . The non-linear transformation $f^{\text{NL}}(\cdot)$ is an essential component for the representation power of neural networks [117]. Formerly, the selected activation functions were either the sigmoidal activation function (3.9a) or the hyperbolic tangent activation function (3.9b), due to their mathematical properties, i.e., they are differentiable, monotonically increasing, and bounded (see Fig. 3.5). Their mathematical formulations are expressed as follows:

$$f^{\text{sig}} = \frac{1}{1 + e^{-x}} \quad (3.9a)$$

$$f^{\text{tanh}} = \frac{e^x - e^{-x}}{e^x + e^{-x}} \quad (3.9b)$$

With the advent of deeper FFNN, today's standard choice of the activation function is rather the rectified linear unit (ReLU) (3.10a), for which experimental results suggest that it improves the convergence of deeper models with faster computation properties. However, due to its derivative behavior (see Fig. 3.5), the ReLU unit may be stuck to output a zero value for any training sample, thereby no longer contributing to the network output (which is called a 'dying ReLU'). For overcoming this issue, other activation functions were designed with a small slope for negative values such as the exponential linear unit (ELU) (3.10b). Let γ an hyperparameter (typically equals to 1), the ReLU and ELU activation functions are expressed as:

$$f^{\text{ReLU}} = \max(x, 0) \quad (3.10a)$$

$$f^{\text{ELU}} = \begin{cases} x & x > 0 \\ \gamma(e^x - 1) & x \leq 0 \end{cases} \quad (3.10b)$$

To be enclosed in the non-linearity region of the activation functions, the data are usually scaled between $[-1, 1]$ before entering the neural network. This procedure does not degrade the information contained in the data, and is necessary to fully leverage the non-linearity ability of neural models. In this work, the input \mathcal{X} and output \mathcal{Y} data (that are continuous) are rescaled using a standardization approach:

$$\hat{z} = \kappa_{\min} + \frac{(z - z_{\min})(\kappa_{\max} - \kappa_{\min})}{z_{\max} - z_{\min}} \quad (3.11)$$

where \hat{z} , z_{\min} and z_{\max} are, respectively, the standardized, maximum and minimum values of a time series \mathcal{Z} , and the values κ_{\min} , κ_{\max} determine the range of

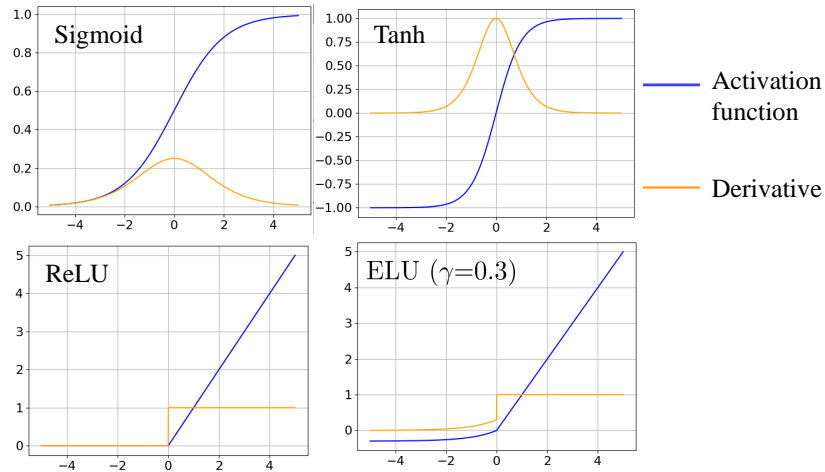


Figure 3.5.: The different activation functions and their derivative.

the scaled values. Note that the data can also be normalized so as to have a zero mean and a standard deviation equal to one. Both scaling procedures can be used, and, in our experiments, do not lead to significant differences in the results.

In addition to continuous inputs, time series forecasting can be feeded with calendar-based features, which help the model to better capture seasonal patterns. For instance, the system imbalance is strongly related to human activity, which thus results in daily and weekly periodicities. By nature, calendar-based features are categorical variables, e.g., the hour of the day $h = \{0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23\}$, for which both rescaling procedures (i.e., standardization and normalization) are not adequate for handling their discontinuities. Indeed, the relative importance between time data is not easily quantified by a numerical value. For instance, the second hour of the day is not 2 times more important than the first one. In this context, a first approach that may be envisaged is one-hot encoding, which consists in representing the categorical time data in a mutually exclusive binary representation [118]. Concretely, the h -value of the hour of the day is mapped to a vector of 24 binary inputs $\mathbf{M}^h \in \mathbb{R}^{24}$, one for each hour of the day, where only one digit can be equal to 1. For example, $h = 2$ is one-hot encoded as $\mathbf{M}^2 = [0 \ 0 \ 1 \ 0]$. One issue with this approach is that it may increase drastically the dimensionality of the input space (e.g., a vector of 96 binary variables encoding all quarter hour of the day). For overcoming this issue, embedding layers can be used, which linearly transforms the one-hot encoded calendar-based feature into a fixed-dimensional vector [119]. Considering the hour of the day, the embedding layer is written

as follows:

$$\mathbf{H} = \mathbf{W}_{M^h}^{\text{emb.}} \cdot \mathbf{M}^h \quad (3.12)$$

where $\mathbf{H} \in \mathbb{R}^D$ is the embedded D -dimensional hour of the day, $\mathbf{W}_{M^h}^{\text{emb.}} \in \mathbb{R}^{D \times 24}$ is the weight parameters to be optimized, and $\mathbf{M}^h \in \mathbb{R}^{24}$ denotes the one-hot encoded hour of the day. Note that the weight parameters of the embedding layer is optimized along with the other neural layers.

The states and outputs of the FFNN solely depend on the provided inputs, with no recurrent connections about previous states of the network. This tends to rapidly increase the dimension of the neural network, and thus, the number of parameters, as the entire input vector must be presented at one go. This could constrain the length of the look-back and look-ahead time windows in time series forecasting. For modeling dynamic processes, recurrent neural networks (RNNs) have been proposed for propagating relevant information within their architecture, without explicitly entering the entire input vector each time the model returns an outcome. The RNNs are presented in the next Section.

3.4. Recurrent Neural Networks

Recurrent neural networks (RNN) are self-connected units, whose recurrent connection h_τ allows the network to have a memory of previous time steps. Given an input sequence $\mathcal{X} = (\mathcal{X}_{t_0+1}, \dots, \mathcal{X}_{t_0+\tau_{\max}})^\top$, the predictions are obtained by iterating the following equations $\forall \tau \in \{t_0 + 1, \dots, t_0 + \tau_{\max}\}$:

$$h_\tau = \mathcal{H}(W_1^{\text{RNN}} \mathcal{X}_\tau + W_2^{\text{RNN}} h_{\tau-1} + b_1^{\text{RNN}}) \quad (3.13a)$$

$$\hat{\mathcal{Y}}_\tau = W_3^{\text{RNN}} h_\tau + b_2^{\text{RNN}} \quad (3.13b)$$

where $W_1^{\text{RNN}} \in \mathbb{R}^{D \times |\mathcal{X}_\tau|}$, $W_2^{\text{RNN}} \in \mathbb{R}^{D \times D}$, $W_3^{\text{RNN}} \in \mathbb{R}^{|\mathcal{Q}| \times D}$, $b_1^{\text{RNN}} \in \mathbb{R}^D$ and $b_2^{\text{RNN}} \in \mathbb{R}^{|\mathcal{Q}|}$ are parameters of the forecasting function. \mathcal{H} is the hidden layer function, which was formerly an element-wise application of a sigmoid function. Note that the hidden dimension D is also a hyperparameter to be tuned.

However, when using an element-wise sigmoid function as \mathcal{H} , the RNNs can suffer issues when backpropagating the gradients for optimizing the weight parameters [120]. Indeed, the gradients may explode or vanish when they are backpropagated along the time dimension (which can represent a large number of time steps), which perturbs the optimization procedure of the forecasting function. To mitigate the latter issue, Hochreiter and Schmidhuber in [121] suggested the long short term memory (LSTM) architecture (shown in Fig. 3.6).

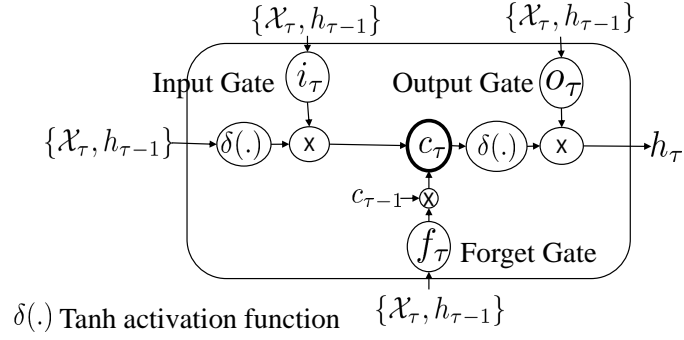


Figure 3.6.: The LSTM unit

The $\mathcal{H}^{\text{LSTM}}$ is composed of a memory cell, whose interactions with the inputs are controlled via three multiplicative gates: the input, forget and output gates. For a time step τ , the LSTM mathematical expressions (with biases omitted) are expressed as :

$$i_\tau = f^{\text{sig}}(W_1^{\text{LSTM}} X_\tau + W_2^{\text{LSTM}} h_{\tau-1}) \quad (3.14a)$$

$$f_\tau = f^{\text{sig}}(W_3^{\text{LSTM}} X_\tau + W_4^{\text{LSTM}} h_{\tau-1}) \quad (3.14b)$$

$$o_\tau = f^{\text{sig}}(W_5^{\text{LSTM}} X_\tau + W_6^{\text{LSTM}} h_{\tau-1}) \quad (3.14c)$$

$$c_\tau = f_\tau c_{\tau-1} + i_\tau f^{\text{tanh}}(W_7^{\text{LSTM}} X_\tau + W_8^{\text{LSTM}} h_{\tau-1}) \quad (3.14d)$$

$$h_\tau = o_\tau f^{\text{tanh}}(c_\tau) \quad (3.14e)$$

where $W_{\{1,3,5,7\}}^{\text{LSTM}} \in \mathbb{R}^{D \times |X_\tau|}$, and $W_{\{2,4,6,8\}}^{\text{LSTM}} \in \mathbb{R}^{D \times D}$ are the parameters of the LSTM forecasting function.

The multiplicative gates i_τ , f_τ , and o_τ allow the memory cell c_τ to store and access information as time elapses, while preserving the gradient flow. For instance, the input gate i_τ can output a value close to 0 such that the corresponding input X_τ will not overwrite the memory cell. Similarly, the forget gate can discard any irrelevant information of previous time steps from the memory cell, while the output gate controls the influence of the memory cell on the output h_τ . Hence, such gating mechanisms allow different LSTM units to capture different time dependencies, which enable the LSTM network to model time series with multiple time scales.

As mentioned, the LSTM allows an access to previously fed information, but no connection with later inputs are permitted in its original form. The bidirectional LSTM (denoted BLSTM, and illustrated in Fig. 3.7) typically enables this access on available future information by processing the input

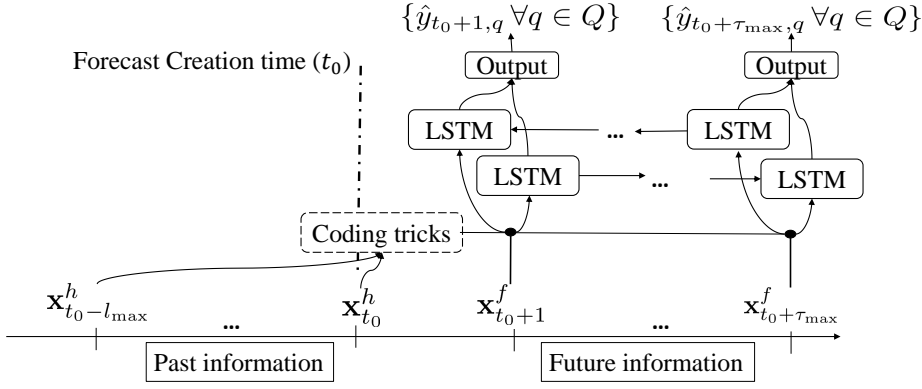


Figure 3.7.: BLSTM neural network processing the information in both positive and negative time directions.

sequence in both directions [122]. Practically, at each time step τ of the prediction horizon, the BLSTM model has access to two hidden states, i.e., \vec{h}_τ that provides a representation of previous events, and \overleftarrow{h}_τ that summarizes the information of the following time steps:

$$\vec{h}_\tau = \mathcal{H}^{\text{LSTM}}(\mathcal{X}_{j,\tau}, \vec{h}_{\tau-1}) \quad (3.15a)$$

$$\overleftarrow{h}_\tau = \mathcal{H}^{\text{LSTM}}(\mathcal{X}_{j,\tau}, \overleftarrow{h}_{\tau+1}) \quad (3.15b)$$

$$\hat{y}_\tau = W_1^{\text{BLSTM}} \vec{h}_\tau + W_2^{\text{BLSTM}} \overleftarrow{h}_\tau + b_1^{\text{BLSTM}} \quad (3.15c)$$

where $W_1^{\text{BLSTM}} \in \mathbb{R}^{|Q| \times D}$, $W_2^{\text{BLSTM}} \in \mathbb{R}^{|Q| \times D}$, and $b_1^{\text{BLSTM}} \in \mathbb{R}^{|Q|}$ are the parameters of the BLSTM output layer.

However, even the BLSTM network is designed to process a fixed-size input vector at each time step, which renders this architecture not flexible enough for encoding both past and future information into the model. As showcased in Fig. 3.7, in a first approach, this can be achieved by a coding trick, where past information are repeated at each time step $\tau \in \{t_0 + 1, t_0 + \tau_{\max}\}$ of the forecasting horizon. However, such coding tricks inevitably augment the size of inputs, and thus, the dimension of the network, which may render the former advantage of recurrent architectures over feed-forward architectures less distinct.

3.5. Bridging Past and Future Input Variables

Efficiently processing the past observed values and the known information about the future is non-trivial with recurrent neural architectures due to the necessity to feed models with a fixed-dimensional input vector. Indeed, the number of past covariates is likely to differ from the number of known inputs about the future. A first successful solution is provided by sequence-to-sequence (also referred to as encoder-decoder) models, which are composed of two different computing blocks for processing respectively past observed and future known information. These architectures have been formerly developed for improving the performance of neural networks in translation applications or speech recognition tasks [123].

The Sequence-to-Sequence Model

The sequence-to-sequence model is represented in Fig. 3.8, where two different blocks process the past observed and future known information. The encoder processes past input data, which is already observed at the forecast creation time t_0 over the look-back window l_{\max} . The aim of the encoder is to map this input sequence into a fixed-length vector c^{enc} that captures all the relevant past dynamics. Based on this past summary, the decoder then generates the multi-horizon predictions along with the known information about the future. In this Chapter, the encoder and decoder blocks are represented by LSTM layers due to their intrinsic capabilities of capturing temporal dynamics. For clarity, the hidden states associated with the encoder blocks are denoted by h_t^{enc} , while h_τ^{dec} is used for the hidden states of the decoder. The encoder block processing the past time steps $t \in \{t_0 - l_{\max}, \dots, t_0\}$ is written as:

$$h_t^{\text{enc}} = \mathcal{H}^{\text{LSTM}}(\mathbf{x}_t^h, h_{t-1}^{\text{enc}}), \quad \forall t \in \{t_0 - l_{\max}, \dots, t_0\} \quad (3.16a)$$

$$c^{\text{enc}} = h_{t_0}^{\text{enc}} \quad (3.16b)$$

Then, the decoder reads this past summary c^{enc} along with the future known inputs \mathbf{x}_τ^f over the future time steps $\tau \in \{t_0 + 1, \dots, t_0 + \tau_{\max}\}$ as follows:

$$h_{t_0}^{\text{dec}} = c^{\text{enc}} \quad (3.17a)$$

$$h_\tau^{\text{dec}} = \mathcal{H}^{\text{LSTM}}(h_{\tau-1}^{\text{dec}}, \{\mathbf{x}_\tau^f; c^{\text{enc}}\}), \quad \forall \tau \in \{t_0 + 1, \dots, t_0 + \tau_{\max}\} \quad (3.17b)$$

$$\mathcal{Y}_\tau = W_1^{\text{dec}} h_\tau^{\text{dec}} + b_1^{\text{dec}}, \quad \forall \tau \in \{t_0 + 1, \dots, t_0 + \tau_{\max}\} \quad (3.17c)$$

where $W_1^{\text{dec}} \in \mathbb{R}^{|Q| \times D}$, and $b_1^{\text{dec}} \in \mathbb{R}^{|Q|}$ are the parameters of the sequence-to-sequence output layer.

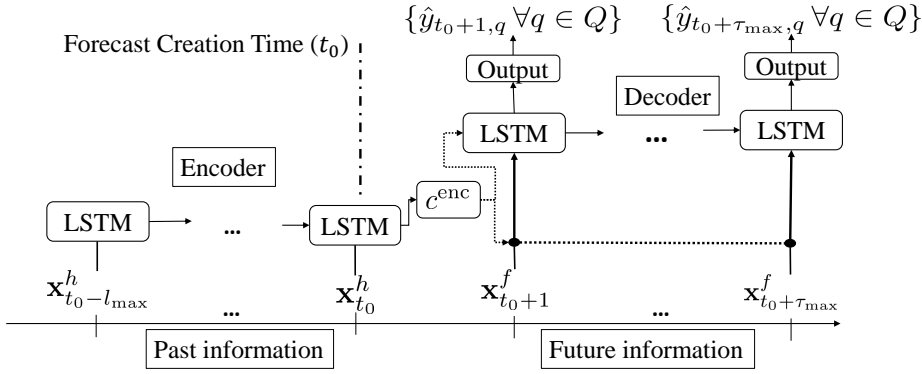


Figure 3.8.: Sequence-to-sequence model for processing both past observed and future known information.

However, the use of a fixed-length vector c^{enc} can be a bottleneck for exchanging information between the encoder and decoder blocks. Indeed, this makes it difficult for the model to cope with large look-back windows. To address this issue, the attention mechanism was proposed, where distinct past vector c_{τ}^{enc} are computed at each time step τ of the forecasting horizon.

The Attention-Based Sequence-to-Sequence Model

The attention mechanism adds a computing layer on top of the encoder block, wherein the input information is mapped into a sequence of vectors $\mathcal{C}^{\text{enc}} = \{c_{t_0+1}^{\text{enc}}, \dots, c_{t_0+\tau_{\text{max}}}^{\text{enc}}\}$ instead of the fixed-length vector c^{enc} . The goal is to selectively adapt the relevant encoder-side information at each time step of the prediction horizon $\tau \in \{t_0 + 1, \dots, t_0 + \tau_{\text{max}}\}$.

Different architectures have been developed to derive this sequence of context vectors \mathcal{C}^{enc} . In this Chapter, we use an attention mechanism inspired by the Bahdanau layer [124]. In that framework, all the hidden states of the encoder are jointly used for constructing each element of the τ_{max} -dimensional context vector \mathcal{C}^{enc} , thus facilitating the representation of long-term dependencies. However, in its traditional form, this attention mechanism only exploits the past information $\{\mathbf{x}_t^h, \forall t \in \{t_0 - l_{\text{max}}, \dots, t_0\}\}$. Here, we extend the attention mechanism on both past and known future data, giving rise to the attention-based sequence-to-sequence model described in Fig. 3.9. The resulting architecture is able to learn complex alignments between the different time steps of the studied horizon. Two attention layers are implemented, and respectively interact with two different encoder blocks. One encoder

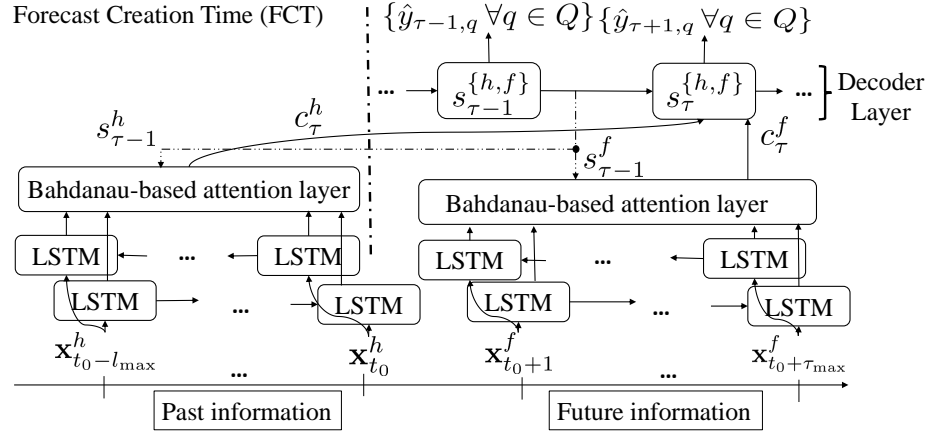


Figure 3.9.: Attention-based sequence-to-sequence model for processing both past observed and future known information.

processes the past data \mathbf{x}_t^h over the look-back steps $t \in \{t_0 - k, \dots, t_0\}$, and the other encoder processes the future data \mathbf{x}_τ^f over the look-ahead steps $\tau \in \{t_0 + 1, \dots, t_0 + \tau_{\max}\}$. The resulting computed information c_τ^h and c_τ^f are then jointly treated into a decoder layer that provides the predictions over the horizon $\tau \in \{t_0 + 1, \dots, t_0 + \tau_{\max}\}$. Note that such a framework allows incorporating longer future contextual information, without being limited to the length of the forecast horizon.

BLSTM models are used as encoders for treating the raw input data in both positive and negative time directions, thus capturing local forward and backward time dependencies. The hidden state \bar{h}_t is defined as the concatenation of the forward and backward BLSTM states, i.e., $\bar{h}_t = \begin{bmatrix} \vec{h}_t \\ \overleftarrow{h}_t \end{bmatrix}$. It thus contains a summary of both preceding and following time horizons, but with a focus on the information at time step t . The hidden states of the two BLSTM models are then used for generating the sequence of context vectors \mathcal{C}^h and \mathcal{C}^f . Without loss of generality, the encoder block processing the past time steps $t \in \{t_0 - l_{\max}, \dots, t_0\}$ is expressed as:

$$\vec{h}_t^h = \mathcal{H}^{\text{LSTM}}(x_t^h, \vec{h}_{t-1}^h), \overleftarrow{h}_t^h = \mathcal{H}_{\text{LSTM}}(x_t^h, \overleftarrow{h}_{t+1}^h) \quad (3.18a)$$

$$\mathcal{C}^h = f^{\text{att}}\left(\left\{\overleftarrow{h}_{t_0-l_{\max}}^h, \dots, \overleftarrow{h}_{t_0}^h\right\}\right) \quad (3.18b)$$

where $f^{\text{att}}(\cdot)$ is the user-defined attention mechanism. The bahdanau-based

attention mechanism is used, which is detailed as follows:

$$c_\tau^h = \sum_{t=t_0-l_{\max}}^{t_0} \alpha_{\tau t} \bar{h}_t^h, \quad \forall \tau \in \{t_0 + 1, \dots, t_0 + \tau_{\max}\} \quad (3.19)$$

where the alignment value $\alpha_{\tau t}$ is computed by:

$$\alpha_{\tau t} = \frac{e^{r_{\tau t}}}{\sum_{t=t_0-l_{\max}}^{t_0} e^{r_{\tau t}}} \quad (3.20)$$

with $r_{\tau t}$ aims at quantifying the degree of relevance of the encoder state \bar{h}_t^h at time t for the prediction outcome at time τ . This alignment value is computed based on the encoder state \bar{h}_t^h and the hidden state $s_{\tau-1}^h$ of the decoder layer at time $\tau - 1$ as follows:

$$r_{\tau t} = v^{\text{att}} f^{\text{tanh}}(W_1^{\text{att}} s_{\tau-1}^h + W_2^{\text{att}} \bar{h}_t^h) \quad (3.21)$$

where $v^{\text{att}} \in \mathbb{R}^{1 \times D}$, $W_2^{\text{att}} \in \mathbb{R}^{D \times D}$ and $W_1^{\text{att}} \in \mathbb{R}^{D \times D}$ are parameters. The hyperparameter D is the dimension of the concatenated hidden state of the BLSTM network. Note that a similar alignment architecture is used to obtain the context sequence \mathcal{C}^f .

Then, the outputs c_τ^h and c_τ^f of both attentional layers are then processed by two LSTM decoders of dimension D , which respectively yield the following hidden states s_τ^h and s_τ^f :

$$s_\tau^h = \mathcal{H}^{\text{LSTM}}(c_\tau^h, s_{\tau-1}^h), s_\tau^f = \mathcal{H}^{\text{LSTM}}(c_\tau^f, s_{\tau-1}^f) \quad (3.22)$$

Finally, the outputs $[s_\tau^h; s_\tau^f]$ of both LSTM decoders (3.22) are concatenated, and are linearly transformed to generate the predictions $\{\hat{y}_{t_0+1, \forall q \in \mathcal{Q}}, \dots, \hat{y}_{t_0+\tau_{\max}, \forall q \in \mathcal{Q}}\}$.

3.6. Training and Inference of Neural Networks

The training of neural networks aims at finding the optimal weight parameters by minimizing an error function \mathcal{L} . This optimization process is traditionally performed using gradient backpropagation [70]. In short, gradient backpropagation is an algorithm that can be summarized in four stages:

1. The forward pass, which consist in outputting the predictions given the current weights and biases values;

2. The computation of an error signal, which is obtained by comparing the targeted values with the predictions using the loss function \mathcal{L} ;
3. The backward pass, where each partial derivatives of the loss function with respect to each weights and biases of the neural model are computed;
4. The updates of weights and biases, which is based on the gradient descent method.

This learning process is iterated by passing several times throughout the whole historical dataset until a convergence criterion is achieved. The gradient descent is used for adjusting the network weights in the direction of the negative error gradient of the loss function \mathcal{L} . The first-order gradient descent method (in its simplest form) is written as:

$$w_{i,j}^{k+1} = w_{i,j}^k - \alpha \frac{\partial \mathcal{L}}{\partial w_{i,j}^k} \quad (3.23)$$

where $w_{i,j}^k$ is the weight between the processing units i and j of two successive layers at iteration k of the backpropagation algorithm. The learning rate $\alpha \in [0, 1]$ scales the step size in the direction of the negative error gradient of the loss function \mathcal{L} . In this report, the Adam approach [125] is used as gradient descent method, which automatically and individually adapts the learning rate α for each network parameter in order to escape local optima during the training phase.

The weight updates can be done in a batch or online mode. In batch mode, the weights and biases of the neural model are only updated after computing the loss function over the whole historical dataset. Regarding the online mode, the weights and biases are iteratively updated based on the loss function of each sample of the historical dataset. The batch mode allows a faster convergence of the optimization procedure, while the stochasticity introduced by the online mode can help escaping from local minima. In this report, a mini-batch mode is preferred, which consists in updating the weight parameters based on the loss function of small groups of samples (i.e., 96 samples representing a daily sequence in our case study), thereby providing a compromise between both approaches.

The computation of the value $\frac{\partial \mathcal{L}}{\partial w_{i,j}^k}$ in (3.24) can rapidly become very complex and computationally intensive for weights and biases located on first layers. For decomposing this mathematical complexity, the backward pass in stage 3. recursively applies the chain rule algorithm. For instance, recalling the notations from Eq. (3.8) of the FFNN in Section 3.3, the partial derivative of

the weight parameters $w_{\tau_q d}^{\text{FFNN}}$ and w_{dx}^{FFNN} are expressed as:

$$\frac{\partial \mathcal{L}}{\partial w_{\tau_q d}^{\text{FFNN}}} = \frac{\partial \mathcal{L}}{\partial \hat{y}_{\tau, q}} \cdot \frac{\partial \hat{y}_{\tau, q}}{\partial a_{\tau_q}} (= 1) \cdot \frac{\partial a_{\tau_q}}{\partial w_{\tau_q d}^{\text{FFNN}}} \quad (3.24a)$$

$$\frac{\partial \mathcal{L}}{\partial w_{dx}^{\text{FFNN}}} = \frac{\partial \mathcal{L}}{\partial \hat{y}_{\tau, q}} \cdot \frac{\partial \hat{y}_{\tau, q}}{\partial a_{\tau_q}} (= 1) \cdot \frac{\partial a_{\tau_q}}{\partial b_d} \cdot \frac{\partial b_d}{\partial a_d} \cdot \frac{\partial a_d}{\partial w_{dx}^{\text{FFNN}}} \quad (3.24b)$$

Note that the chain rule is even more expanded when optimizing recurrent neural networks due to the recurrent connection over the time dimension. In this case, the overall optimization process is called backpropagation through time. Overall, the difficulty of coding the backward pass has been greatly reduced with the introduction of libraries such as Tensorflow [126] written in the open source language Python, as routines are included to automatically differentiate any neural network architectures.

The backpropagation algorithm ends when a convergence criterion is reached. For maximizing the generalization capability to unseen data, the commonly adopted criterion is early stopping. This criterion necessitates to divide the whole historical dataset into two separate sets, i.e., the training and validation sets. The weight parameters of the model are updated using the training set, while the validation set is used for assessing the generalization abilities of the model (i.e., no optimization is done on the signal errors when travelling the validation set). The optimization process is stopped when no performance improvement is apparent on the validation set. Indeed, as soon as the error on the validation set stagnates or increases, this signifies that the network starts to overfit, i.e., exploiting artifacts in the data (e.g., erratic error measurements) rather than relevant patterns. A typical illustration of the evolution of training and validation curves for a FFNN is showcased in Fig. 3.10, where the weight parameters of the final model are selected based on the minimum error of the validation set.

The dimension of the network also influences the generalization capabilities of the model. Indeed, a model too simplistic will not be able to replicate the underlying data generation process, while a model excessively complex is more prone to the overfitting issue. The dimension of the network is described by two hyperparameters, i.e., the number of hidden layers L and the number of processing units D within each layer. These two hyperparameters are task-dependent parameters, whose optimal values are not known beforehand. Hence, in order to find the optimal model complexity during the training phase, different configurations need to be tested, which requires each time a complete training of a new model from scratch. In this report, the selection of




Figure 3.10.: Evolution of the training and validation curves for a FFNN with 50 processing units.

hyperparameters is achieved by random search, which uniformly draws trials from each configuration space of hyperparameters (e.g., $L = \{1, 3, 5\}$ and $D = \{5, 10, 50, 100, 250\}$) [127]. The final configuration is the one achieving the lower validation loss.

In our case studies, the loss function \mathcal{L} is the pinball loss function. However, in case of perfect prediction, the pinball loss cannot be differentiated because of the kink (sharp corner point) of this function at this particular point. To address this issue, a smooth approximation of the pinball loss is constructed by including the Huber norm within the loss function [128]. The idea is to replace the L_1 norm by the (continuously differentiable) L_2 norm when the error is lower than a (preferably small) user-defined threshold ϵ (here, we arbitrarily use $\epsilon = 10^{-6}$):

$$H(y_\tau, \hat{y}_{\tau,q}) = \begin{cases} \frac{(\hat{y}_{\tau,q} - y_\tau)^2}{2\epsilon} & |\hat{y}_{\tau,q} - y_\tau| \leq \epsilon \\ |\hat{y}_{\tau,q} - y_\tau| - \frac{\epsilon}{2} & |\hat{y}_{\tau,q} - y_\tau| > \epsilon \end{cases} \quad (3.25)$$

where the values $y_{\tau,q}$ are the outputs of the forecaster, and y_τ the actual (ground truth) observations. The approximated pinball loss E_Q can then be calculated as:

$$E_Q = \sum_{\tau \in T} \sum_{q \in Q} \begin{cases} q \cdot H(y_\tau, \hat{y}_{\tau,q}) & \hat{y}_{\tau,q} < y_\tau \\ (1 - q) \cdot H(y_\tau, \hat{y}_{\tau,q}) & \hat{y}_{\tau,q} \geq y_\tau \end{cases} \quad (3.26)$$

As the resulting loss function is differentiable, the neural networks can thus be trained using the gradient-based backpropagation method. Note that quantile crossing issue may arise when predicting a discrete set of quantiles. In this report, a naive rearrangement of the predicted q -quantiles is conducted in ex-post, i.e., the q -quantiles are sorted in ascending order at each time step of the prediction horizon after they are predicted [129]. Note that this procedure is also performed for the benchmark methods.

Once trained, the neural models are then used on novel input conditions, i.e., the test set, for inferring their probabilistic predictions. The inference process is simply the stand-alone computation of the forward pass (stage 2. of the optimization procedure). Regarding computational times between training and inference stages, for instance, an attention-based sequence-to-sequence model with 131k parameters takes around 1 minute for being trained over one pass of 103872 training samples using a mini-batch size of 96, while its inference time for generating its predictions on a sample of the test set is lower than one second. This difference between training and inference time also stands for other time series forecasting approaches (e.g., tree-based or econometric models), which renders them operational for close-to-real-time purposes.

3.7. Benchmark

We compare the proposed neural network architectures with a wide range of techniques for the probabilistic forecasting of system imbalances. First, two naive methodologies are implemented:

- A step-wise averaging model (Step-Avg), where the system imbalance distribution of each forecasting time step is computed based on the average of all past observations corresponding to this specific period of the day.
- A probabilistic generalization of persistence (Persistence) based on a random walk model. The forecast assumes a Gaussian distribution where the mean is given by the last available system imbalance realization, and the variance is determined by exponential smoothing of previous squared errors [130].

Then, three other well-established models in time series forecasting are also implemented:

- An Auto-Regressive Moving Average (ARMA) model, for which we compute prediction intervals assuming that the residuals are uncorrelated and normally distributed [131].

- A quantile regression forest (QRF), i.e., a bagging-based ensemble method, in which the outcomes of independent regression trees are used for estimating the conditional distribution [132]. The forest is set to 500 trees. Note that the random forest method is more detailed in Subsection 5.2.3 of Chapter 5 for the interested reader.
- A gradient boosting regression tree (QGBRT) trained with the quantile loss. New regression trees are sequentially created to predict the residuals of the previously generated ones [133]. The number of boosting stages is fixed to 100 with an early stopping criterion. Note that the gradient boosting regression tree method is more detailed in Subsection 5.2.3 of Chapter 5 for the interested reader.

It should be noted that the ARMA model is only fed with past system imbalance observations, while tree-based ensemble models (QRF and QGBRT) have access to the same input data as the neural models. Note that, to avoid tree-based models to have an output of dimension $\tau_{\max} \cdot |Q|$, a different model is trained for each time step τ of the prediction horizon. This reduces the output size of each model (thus facilitating training), but at the expense of an increased training time. In particular, each of the $\tau_{\max} \cdot |Q|$ QGBRT models takes around 3 minutes to train (for a total time of more than 45 minutes), while the whole training time of the τ_{\max} QRF models takes 73 minutes. Finally, the neural models are denoted as follows:

- FFNN stands for the 1-layer feed-forward neural network, while S-FFNN represents the stacked feed-forward neural network.
- LSTM and BLSTM denotes respectively the Long Short Term Memory recurrent neural network and Bidirectional Long Short Term Memory recurrent neural network.
- Seq2Seq and Att-Seq2Seq are the sequence-to-sequence model and attention-based sequence-to-sequence model presented.

For each forecaster (except the parameter-free naive approaches), the optimal model complexity is identified using the random search procedure. The same budget of trials is used across all benchmarks. For illustration purposes, the search ranges for hyper-parameters of the Att-Seq2Seq is listed hereafter:

- Dimension of the network $D = \{6, 12, 24, 32, 64, 128\}$;
- number of past time steps $l_{\max} = \{4, 8, 12, 16, 20, 32, 48\}$;
- initial learning rate $a = \{10^{-2}, 10^{-3}, 10^{-4}\}$.

3.8. Metrics for Probabilistic Forecasting

When dealing with probabilistic forecasting, two (potentially conflicting) notions are important to consider, i.e. calibration (or reliability) and sharpness [134]. Calibration refers to the statistical correctness between the predicted quantiles and the true distribution (i.e., actual observations). Sharpness is a simple measure of the concentration (width) of the predictive distributions. Powerful forecasts must thus find a trade-off between maximizing the sharpness (concentrated intervals), while ensuring that the reliability of the predicted intervals is preserved. To comprehensively assess this compromise between reliability and sharpness, three probabilistic scoring tools are employed.

Firstly, we use the pinball loss E^Q weighted across all q -quantiles of interest (in this work, $Q = \{0.05, 0.15, 0.25, 0.35, 0.45, 0.5, 0.55, 0.65, 0.75, 0.85, 0.95\}$).

$$E_\tau^Q = \sum_{q \in Q} q \max(0, y_\tau - \hat{y}_{\tau,q}) + (1 - q) \max(0, \hat{y}_{\tau,q} - y_\tau) \quad (3.27)$$

where $y_{\tau,q}$ are the forecasted quantiles, and y_τ the actual observations of the system imbalance (in MW). A lower E_τ^Q score indicates a better probabilistic forecast.

However, by scoring all q -quantiles in one final metric, the pinball loss may hide low reliability levels for extreme quantiles. For instance, even if the 5% quantile forecasts completely fail, it may have a very limited impact on the total score. To address this issue, the Winkler score is implemented, which quantifies the forecast quality for different prediction intervals. For a prediction interval of $(1 - \beta)100\%$, the Winkler score E_τ^W is defined as:

$$E_\tau^W = \begin{cases} \epsilon_\tau, & L_\tau \leq y_\tau \leq U_\tau, \\ \epsilon_\tau + 2(L_\tau - y_\tau)/\beta, & y_\tau < L_\tau, \\ \epsilon_\tau + 2(y_\tau - U_\tau)/\beta, & y_\tau > U_\tau, \end{cases} \quad (3.28)$$

where $L_\tau = \hat{y}_{\tau,\beta/2}$ and $U_\tau = \hat{y}_{\tau,1-\beta/2}$ are the lower and upper bounds of the prediction interval defined by the confidence level β , and $\epsilon_\tau = U_\tau - L_\tau$ is the interval width. In this report, the Winkler score E_τ^W is calculated for $\beta = \{0.1, 0.5, 0.9\}$.

If an actual realization y_τ is within the predicted interval $[L_\tau, U_\tau]$, the Winkler score E_τ^W is a direct measure of sharpness. Otherwise, a penalty term, whose value depends on the severity of the forecast error, is added for reflecting the deficiency in reliability. Finally, we also use the continuous ranked probability

score (CRPS), defined as:

$$E_{\tau}^{\text{CRPS}} = \int_u \left(\hat{F}(u) - \theta(u - y_{\tau}) \right)^2 du \quad (3.29)$$

where $\hat{F}(\cdot)$ is the cumulative distribution function (cdf) defined by the predicted quantiles $\hat{y}_{\tau,q}$, and $\theta(\cdot)$ is the Heaviside step function, taking the value 1 for $u \geq y_{\tau}$, and 0 otherwise.

Eq. (3.29) is a quadratic measure of the difference between the cdf and the observation, which is null in case of a perfect probabilistic forecast. It measures both reliability and sharpness, and has the same unit than the variable of interest. As we evaluate non-parametric predictive densities, we can numerically integrate Eq. (3.29) to compute the CRPS score.

3.9. Case study: Probabilistic Forecasting of Belgian System Imbalances

The case study is conducted on an Intel® Core™ i7-3770 CPU @ 3.4 GHz with 16 Gb of RAM. The forecasting methods are implemented using Python 3.6.0 with the Keras library (along with the TensorFlow backend) for the neural models, with the scikit-learn library for the tree-based ensemble models, and with the statsmodels for the econometric models. We focus on the probabilistic predictions of the system imbalance (denoted SI) with a 15-minute time granularity. As mentioned in Section 3.2, the system imbalance is assumed equals to the opposite value of the Net Regulation Volume (NRV), which is activated by the Transmission System Operators (TSO). The data spans from 2016-1-1 to 2019-12-31, for a total of four years of data. Specifically, the first three years of data are used to train and validate the models with a ratio of 85%-15%, while the fourth year is used for testing. Each quarter-hourly step of the database is used as a forecast creation time t_0 . A prediction horizon of 4 hours is selected, which corresponds to $\tau_{\max} = 16$ time steps, and we compute the 5th, 15th, 25th, 35th, 45th, 50th, 55th, 65th, 75th, 85th, 95th percentiles of the target distribution (i.e., $|Q| = 11$) for each of these time periods.

TSOs have the duty to publish a wide range of information for promoting a transparent and non-discriminatory market such as actual measurements (e.g., electricity demand and power production, which are here denoted by the superscript h), day-ahead forecasts of renewable generation and electrical load (denoted here by the superscript f). Additional information, such as the

schedules of conventional generation and merit order proxies of operational balancing prices, may also be provided (also denoted by the superscript f). In the Belgian power system, we have at our disposal $m_h = 14$ historical covariates $\mathbf{x}_{t_0-l_{\max}:t_0}^h$ and $m_f = 15$ known future information $\mathbf{x}_{t_0+1:t_0+\tau_{\max}}^f$. These inputs are gathered as followed:

- the Net Regulation Volume ($\text{NRV}^h \in \mathbb{R}^1$);
- the imbalance price ($\lambda^{h,\text{RT}} \in \mathbb{R}^1$);
- the Marginal Decremental Price (MDP) and the Marginal Incremental Price ($\lambda^{h,\text{bal.}} \in \mathbb{R}^2$);
- the physical cross-border energy flows with France and Netherlands ($\phi^h \in \mathbb{R}^2$);
- the produced and forecasted wind and photovoltaic powers, with their associated installed capacities ($P^{\{h,f\},\text{renew.}} \in \mathbb{R}^4$);
- the produced and scheduled powers of conventional generators ($P^{\{h,f\},\text{conv.}} \in \mathbb{R}^3$), composed of pump-hydro, gas and nuclear units;
- the measured and forecasted electricity demand of the grid ($L^{\{h,f\}} \in \mathbb{R}^1$);
- the day-ahead electricity prices ($\lambda^{f,\text{DA}} \in \mathbb{R}^1$);
- the merit order proxies of operational balancing prices, i.e., the TSO expected prices corresponding to different volumes of activated reserves $\{-600, -300, -100, 100, 300, 600\}$ MW ($\lambda^{f,\text{bal.}} \in \mathbb{R}^6$).

In parallel, calendar information ($\mathbf{x}^{\{h,f\},\text{cal.}} \in \mathbb{R}^6$), i.e., working days, the day of the week, the hour, the quarter hour, the month and the absolute position of the time step, are also available. They are represented by categorical variables, where, e.g., the quarter hours are described by the set $\{0, 1, 2, 3\}$. For neural models, each calendar information ($\mathbf{x}^{\{h,f\},\text{cal.}} \in \mathbb{R}^6$) is processed through an embedding layer, which transforms the calendar variable into a fixed-dimensional vector. This reduces the dimensionality of the input space (e.g., by avoiding the use of 96 binary variables to encode all quarter hours of the day), while providing a new learned meaningful representation able to capture their relative significance. In this way, time steps with similar properties are placed close to each other in the embedding vector, which cannot be achieved with traditional techniques such as one-hot encoding. Overall, the set of historical covariates $\mathbf{x}_{t_0-l_{\max}:t_0}^h$ is composed of $\{\lambda^{h,\text{RT}}, \text{NRV}^h, \lambda^{h,\text{bal.}}, \phi^h, L^h, P^{h,\text{renew.}}, P^{h,\text{conv.}}, \mathbf{x}^{h,\text{cal.}}\}$, while the set of future known information $\mathbf{x}_{t_0+1:t_0+\tau_{\max}}^f$ contains $\{L^f, P^{f,\text{renew.}}, P^{f,\text{conv.}}, \lambda^{f,\text{DA}}, \lambda^{f,\text{bal.}}, \mathbf{x}^{f,\text{cal.}}\}$. The final configurations of the forecasting models are:

- ARMA model, which considers 12 past values and 3 previous values of past errors, i.e., the autoregressive part $p = 12$ and the moving average part $q = 2$.
- QRF model, with a population of 500 trees fully extended. The look-back window is $l_{\max} = 32$.
- QGBRT model, with a learning rate $\alpha = 0.1$, a maximum depth of 8 per tree, and the number of iterations is determined by using early stopping. The look-back window is $l_{\max} = 32$.
- FFNN model, with $D = 24$ processing units. The look-back window is $l_{\max} = 4$.
- S-FFNN model, with three layers of $D = 24$ processing units. The look-back window is $l_{\max} = 4$.
- LSTM model, with $D = 12$ processing units. The look-back window is $l_{\max} = 4$.
- BLSTM model, with $D = 12$ processing units. The look-back window is $l_{\max} = 4$.
- Seq2Seq model, with $D = 64$ processing units. The look-back window is $l_{\max} = 20$.
- Att-Seq2Seq model, with $D = 32$ processing units. The look-back window is $l_{\max} = 8$.

First, for illustrating the quality of the outcomes obtained using the att-Seq2seq model, probabilistic forecasts (over the 4 hours horizon) for March 29, 2020, at 1:00 am, 5:00 am, and 11:00 am, are depicted in Fig. 3.11. The predicted intervals tend to properly embed the actual system imbalance realizations, suggesting that the volatility of the signal is well captured. We can observe some spikes at hourly intervals, i.e., higher imbalances occur in the first quarter hour of each new hour. This observation arises from the fact that European day-ahead markets are characterized by an hourly time resolution, where the electricity is traded in hourly blocks [MWh/h]. This results in a dispatch of power plants with ramping trajectories between consecutive hours that are causing imbalances in the TSO control area. It is worth noting that several theoretical solutions are proposed to tackle this problem, such as trading power trajectories instead of energy blocks [135], or by introducing asynchronous energy blocks [136].

Fig. 3.12 presents the forecast accuracy in terms of quantile loss of the different models as a function of the prediction horizon to analyze up to which

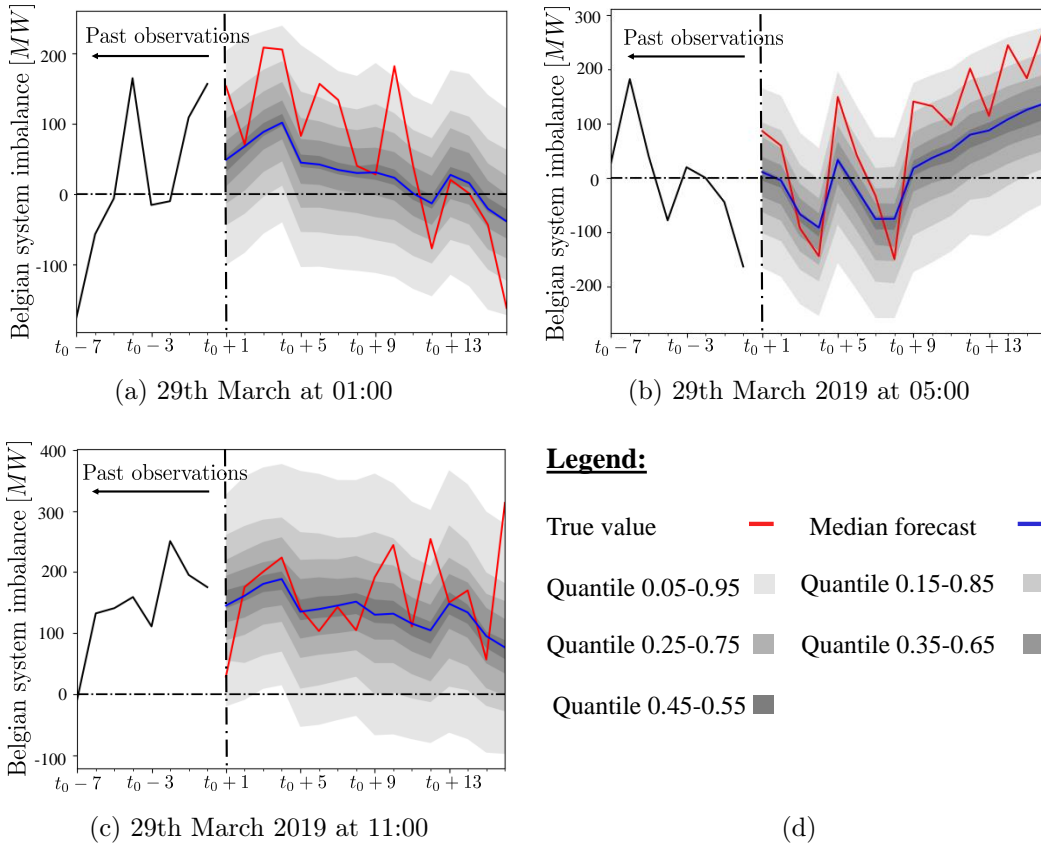




Figure 3.11.: Multi-horizon probabilistic forecasts of the system imbalance on the 29th March 2019 at 01:00 am (Fig. 6.7a), 5:00 am (Fig. 6.7b), and 11:00 am (Fig. 3.11c).



(a)



(b)

Figure 3.12.: Quantile score of different methods (averaged over the whole test set) for the prediction horizon $\{t_0 + 1, \dots, t_0 + \tau_{\max}\}$. Outcomes are split into two sub-figures, i.e., the first one comparing the att-Seq2Seq model with traditional benchmarks, the second one comparing the att-Seq2seq with the other neural models.




Figure 3.13.: Autocorrelation function of the system imbalance signal

time span it is relevant to anticipate system imbalances. Overall, the differences in performance (between all forecasters) quickly decrease over time since all methods tend towards a naive representation of the historical distributions (differentiated between the quarter hour), thus illustrating the difficulty to capture the extreme volatility and unpredictability of the system imbalance over long horizons. In that respect, the naive benchmark (Step-average) estimating the future system imbalance distributions independently for each time step (based on the historical system imbalance values measured each day at the corresponding period) provides a competitive baseline, which is not easily overcome. Due to the high variability of the system imbalance, the Persistence model performs very badly. Indeed, it simply propagates past observations without the ability to infer the most likely future realizations. This behavior can be explained by the autocorrelation function of the system imbalance (Fig. 3.13), which shows moderate time dependencies among consecutive system imbalance realizations. However, we observe peaks at indices 4 and 8 (corresponding to hourly dependencies), which explains the improved accuracy of Persistence at these time steps (with respect to surrounding steps). We also see that the ARMA model achieves poor prediction performance, which may arise from different reasons. First, it only leverages past system imbalance values (thus neglecting the potential of exogenous information) in a linear framework. Second, the prediction intervals of ARMA are computed analytically, assuming that residuals are uncorrelated and normally distributed, which may lead to poor accuracy when these assumptions are violated.

Then, it can be observed in Fig. 3.12 that the att-Seq2Seq model outperforms all (naive and state-of-the-art) benchmarks over the 4-hour prediction horizon. In particular, it yields an average improvement of 1.5% and 3.6% compared to Seq2seq and QGBRT. QRF also provides reliable outcomes,

Models	CRPS [MW]	Winkler score $E_{t_0+1}^W$ [MW]		
		$\beta = 0.1$	$\beta = 0.5$	$\beta = 0.9$
Step-Avg	60	523.2	299.2	205.1
Persistence	65.2	430.7	247.9	167.9
ARMA	59.3	379.6	221.9	152.5
QRF	54.1	389.2	233.2	160.2
QGBRT	54.9	357	208	142.1
FFNN	58	414.8	240.1	163.6
S-FFNN	56.2	373.5	220.4	151.6
LSTM	54.9	366.7	212.2	146
BLSTM	53.8	378.3	220.3	151.4
Seq2Seq	53.7	362.2	211.6	144.7
Att-Seq2Seq	52.9	356.2	210.5	144.2

Table 3.1.: CRPS score averaged over the entire prediction horizon and the Winkler score for the first step of the prediction horizon $t_0 + 1$.

but it is less accurate than its QGBRT counterpart. Although tree-based models have the advantage of processing any type of (continuous and integer) inputs without the need to normalize these data, they still lack the processing abilities of advanced neural architectures. Another interesting result is that the S-FFNN (3 hidden layers) obtains better results than the FFNN, which highlights the importance of depth for leveraging the full potential of neural networks. Moreover, the LSTM and BLSTM have better performances than the S-FFNN, which illustrates the benefit to design tailored architectures for processing the temporal dependencies. Indeed, the att-Seq2seq model, whose temporal nature is designed to optimally exploit the past and future temporal dynamics of inputs, seems to achieve this objective in the lower part of Fig. 3.12 by outperforming all other neural networks.

To complement these results, the CRPS over the entire prediction horizon and the Winkler scores for different reliability levels at time step $t_0 + 1$ are computed for all models, and the results are provided in Table 3.1. We observe similar trends in both metrics. In particular, the Att-Seq2seq model achieves a higher accuracy in average for the CRPS compared to the other forecasting models. However, it should be noted that the QGBRT yields the best results for the first time step $t_0 + 1$ (according to the Winkler score) for the prediction intervals $\beta = \{0.5, 0.9\}$, which emphasizes its prediction ability for narrower intervals. Yet, for extreme quantiles, the Att-Seq2Seq obtains better prediction intervals according to the Winkler scores.

3.10. Conclusion

This Chapter is devoted to the multi-horizon probabilistic forecasting of the system imbalances. The system imbalance is a highly variable and uncertain signal, which directly results from the real-time operating conditions of the system. This information is essential for providing real-time balancing services, as the direction of the system imbalance impacts the imbalance price regime (either low via the Marginal Decremental balancing energy Price (MDP) or high via the Marginal Incremental balancing energy Price (MIP)).

The following neural architectures of gradual complexity are assessed and compared: i) shallow (i.e., 1-layer) and stacked feed-forward neural networks, ii) Long Short Term Memory neural network and their bidirectional counterpart, and iii) the sequence-to-sequence with and without attention mechanisms. All these neural architectures have been detailed, and are sequentially suggested for better capturing the specificities pertaining to time series forecasting (whose inputs can be composed of past observed and future known inputs). The outcomes show that advancements in terms of neural architecture are accompanied with an increase of forecasting performance. Practically, the sequence-to-sequence Long Short Term Memory neural networks augmented with attention mechanisms show the highest accuracy compared to other benchmark outcomes (including econometric models and tree-based ensemble methods). This suggests that neural architectures suited to the temporal nature of inputs allow to generate more accurate (tightened) quantiles in comparison with other neural architectures, which paves the way towards further research in this direction.

The next Chapter is focused on valuing these probabilistic forecasts in decision-support tool dedicated to provide real-time balancing services.

Chapter Publications

- **J. Bottieau**, F. Vallée, Z. De Grève, J-F. Toubeau, "Leveraging Provision of Frequency Regulation Services from Wind Generation by Improving Day-Ahead Predictions using LSTM Neural Networks," in 2018 IEEE International Energy Conference (ENERGYCON), Limassol, Cyprus, 2018.
- J-F. Toubeau, **J. Bottieau**, F. Vallée and Z. De Grève, "Deep Learning-Based Multivariate Probabilistic Forecasting for Short-Term Scheduling in Power Markets," in IEEE Trans. Power Syst., vol. 34, no. 2, pp. 1203-1215, 2019.

- J-F. Toubreau, T. Morstyn, **J. Bottieau**, K. Zheng, D. Apostolopoulou, Z. De Grève, Y. Wang, and F. Vallée, "Capturing Spatio-Temporal Dependencies in the Probabilistic Forecasting of Distribution Locational Marginal Prices," in *IEEE Trans. Smart Grid*, vol. 12, no. 3, pp. 2663-2674, 2021.

CHAPTER 4.

Risk-Aware Stochastic Provision of Real-Time Balancing Services

The real-time balancing services are provided by the intentional imbalance positions of Balance Responsible Parties (BRPs) in the opposite direction of the total net system imbalance. Such services are typically incentivized in a single imbalance pricing scheme, where all BRPs imbalances are settled at a unique price. This Chapter proposes risk-aware stochastic mathematical formulations for providing these services. More particularly, the bi-level methodology is firstly investigated for capturing the interaction between the intentional BRP imbalance position and the clearing of the balancing energy market, which allows avoiding excessive imbalance positions that overstep the system imbalance. Then, for exploiting at best the probabilistic predictions of the system imbalance, the uncertainty around the predicted variable is also modeled in the decision-support tool. This allows improved out-of-balance decisions, giving the possibility for actors to manage the financial risks, i.e., the possibility that an actor's financial outcome deviates adversely from its expectation, associated with their energy positions.

This Chapter is organized as follows. Section 4.1 states the market assumptions for modeling the single price imbalance settlement in the Belgian context, followed by a succinct presentation of the market application. Section 4.2 shows the decision-support tool formulated as a bi-level optimization problem, while Section 4.3 exposes the reformulation steps for solving the bi-level structure by off-the-shelf optimization solvers. Then, two risk-aware stochastic optimization methods, based on the bi-level model, are presented, namely robust optimization and scenario-based stochastic programming. The two optimization methods differ from their uncertainty modeling paradigm (which is detailed in Section 4.4): i) robust optimization in Section 4.5 determines the uncertainty via a deterministic uncertainty set, and ii) stochastic programming in Section

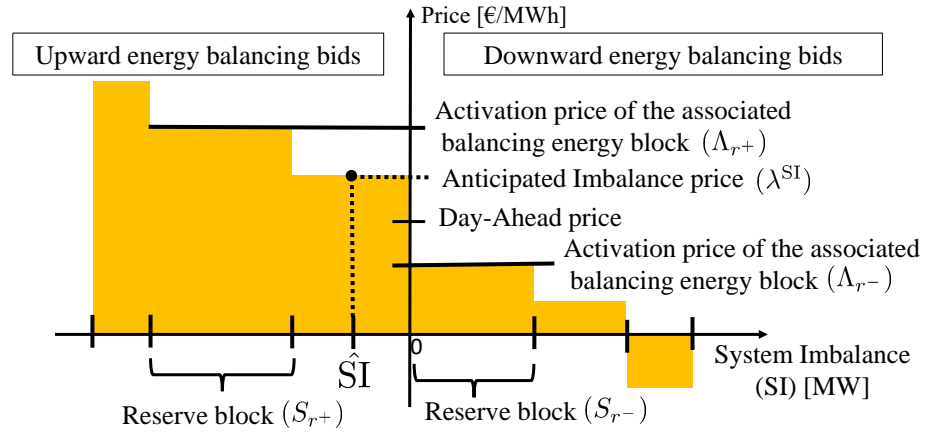


Figure 4.1.: Reconstruction of the balancing energy market from the published Belgian TSO data.

4.6 represents the uncertainty under a set of probability-weighted scenarios. Then, extensive out-of-sample analysis are performed in Sections 4.7-4.9 using real-world market data from the Belgian power systems. More specifically, Section 4.7 analyses the practical gains of generating more accurate probabilistic predictions for providing real-time balancing services through the robust model. Then, Section 4.8 investigates the economic interest of modeling the decision-support tool as a bi-level structure. Finally, Section 4.9 validates and compares different financial risk managements that can be adopted by BRPs for robust optimization and scenario-based stochastic programming.

4.1. Market Assumptions and Application

The single imbalance pricing scheme creates thus an opportunity cost for the Balance Responsible Parties (BRPs) aggravating the system imbalance and an opportunity profit for the BRPs helping the system to be balanced. This mechanism is currently implemented in Netherlands, Germany and Belgium. In a 2011 analysis comparing the Dutch and German imbalance management systems [137] (which were respectively based on single and dual imbalance pricing), the TSO TenneT suggests that the Dutch system has a higher macro-economic efficiency due to the incentives provided to Balance Responsible Parties (BRPs) in correcting imbalances. Besides, after that Germany has adopted the single imbalance pricing scheme, the study in [138] shows that German BRPs take intentional imbalance positions to

financially optimize their own portfolio, while having a positive impact on the system balancing. In that vein, the Belgian Commission of Regulation of Electricity and Gas (CREG) explicitly states in [139] that each BRP is required to contribute to a balanced power system, either by maintaining a balanced portfolio or by holding an imbalanced position in the direction that helps the power system as a whole. This is aligned with their 2017 monitoring report [140], which shows that the Belgian average daily profile of positive and negative system imbalances are rather low in 2017 compared with the period 2007-2013 (where dual imbalance pricing was predominantly implemented). The provision of real-time balancing services (also called passive contribution) is also supported by some market stakeholders at European level, where, for example, the European Federation of Energy Traders (EFET) in [141] argues that the imbalance price should not be regarded as a penalty to force market parties to stick to their schedules. Rather, the imbalance price should reflect the value of electricity in real time, which, as such, provides the correct economic signal to avoid imbalances or help the system.

Hence, the Belgian power system is selected as the candidate for assessing the decisions of the developed decision-support tools, as the Belgian imbalance settlement is perfectly aligned with the European guidelines on the balancing markets. The main assumption concerns the definition of the system imbalance, whose value is given by the opposite value of the net Regulation Volume (NRV) signal. This implicitly assumes that the Transmission System Operator (TSO) has perfectly responded to the actual system imbalance. This definition of system imbalances is aligned with the Belgian imbalance pricing scheme over the years 2016-2019, in which the imbalance price regime was driven by the sign of the NRV signal at each corresponding quarter hourly imbalance settlement period [115].

For promoting efficient and transparent balancing energy markets, the Belgian TSO provides day-ahead information with near real-time updates on the available balancing energy levels ($S_{r+,-}$) and associated activation costs ($\Lambda_{r+,-}$) for upcoming 15-minute imbalance settlement periods [95]. As illustrated in Fig. 4.1, these different activation costs ($\Lambda_{r+,-}$) corresponding to the activation of different balancing energy levels ($S_{r+,-}$) can be used to construct the merit order curves (in yellow in Fig. 4.1), which provides a proxy of the balancing energy market-clearing process. Yet, it should be noted that this proxy implies two additional assumptions w.r.t. the actual Belgian balancing energy market: i) the published balancing energy bids assume an imbalance price constant over intervals of 100 MW, and ii) although these data cover a large part of the balancing energy products, all balancing energy bids

are not incorporated such as balancing bids of pumped hydro storage units (i.e., principally Coo and La Plate Taille storage units in Belgium) [142].

These market data are an important source of information in our market application, which is the decision problem faced by a BRP owning remaining flexibility margins at the beginning of the quarter hour. In this setting, the BRP has the ability to deviate from its current energy position to restore the balance of the power system, thereby allowing to provide balancing services not defined by the standard energy balancing products. The imbalance position of the market actor (denoted hereafter as $e^{\text{imb},+/-}$), i.e., the remaining deviation between its schedule and its real-time physical position, is optimized by anticipating what could be the imbalance price λ^{SI} at the end of the imbalance settlement period based on forecasts of the system imbalance $\hat{\text{SI}}$. This dependency between the prediction of the system imbalance $\hat{\text{SI}}$, the proxy of the balancing energy market, and the anticipated imbalance price λ^{SI} is showcased in Fig. 4.1. This framework is run sequentially 96 times a day at the start of each 15 minutes imbalance settlement period with a horizon of one time step. At each imbalance settlement period, the BRP thus decides on the direction (downward or upward) and volume of the provided real-time balancing service.

4.2. The Bi-level Model

As the volumes of energy exchanged in the balancing energy market are by nature small (typically between -150 MW and 150 MW for the Belgian power system), the intentional imbalance positions of the BRP may therefore have an impact on the balancing energy market clearing. In this line, Fig. 4.2 illustrates two situations where the intentional imbalance position of the BRP oversteps the system imbalance. In case a), the upward real-time balancing service provided by the BRP is greater than the absolute value of the system imbalance, and, consequently, creates a need of downward balancing. In the end, the positive energy surplus of the BRP is remunerated at a very low, even negative, imbalance price, which penalizes the BRP compared to, e.g., the same quantity offered in the day-ahead electricity market. In case b), an analogue reasoning can be made for a real-time balancing service provided in case of a negative system imbalance. Here, the imbalanced BRP will be charged with a much higher imbalance price in comparison with the day-ahead electricity price for its shortage of energy. For alleviating such situations, this interaction between the imbalanced BRP and the clearing of the balancing energy market can be mathematically




Figure 4.2.: a) Excessive imbalance position of a market actor overshooting a negative system imbalance. b) Excessive imbalance position of a market actor overshooting a positive system imbalance.

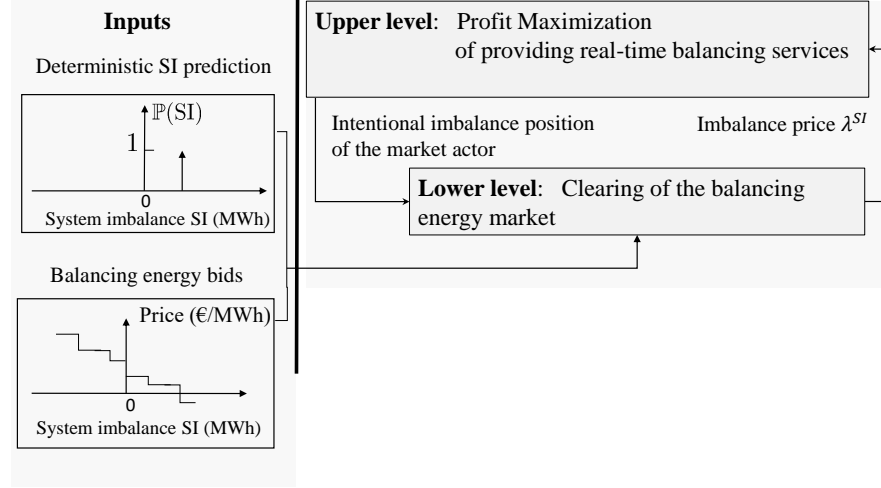


Figure 4.3.: The bi-level decision-support tool in a deterministic framework.

captured by a bi-level formulation, which involves an optimization problem constrained by another optimization problem, see e.g., [78], [92]. Indeed, since the balancing energy market is cleared with the purpose of minimizing the activation costs of the balancing energy products, the description of the clearing mechanism can be conveniently formulated as an optimization problem.

The participation of the actor in the single price imbalance settlement is thus modeled in a bi-level optimization framework, in which the market actor anticipates the clearing of the imbalance settlement scheme (λ^{SI}) by having at its disposal i) a list made available by the TSO of the submitted balancing offers and the associated activation prices, and ii) its own future description of the future system imbalance. The proposed framework is summarized in Fig. 4.3 in a deterministic setting. The upper level (4.1a)-(4.1b) aims at defining the optimal imbalance position ($e^{\text{imb},+/-}$) of the market actor and the lower level (4.1c)-(4.1f) performs the clearing of the balancing energy market. In a deterministic setting, the problem reads as:

$$\max_{\substack{e^{\text{imb},+/-}, \lambda^{SI} \\ \geq 0}} f^{\text{IS}} \left(e^{\text{imb},+/-}, \lambda^{SI} \right) \quad (4.1a)$$

$$\text{s.t. } e^{\text{imb},+/-} \in \Pi_{\text{UL}} \quad (4.1b)$$

$$\min_{\substack{\Theta_{\text{LL}} = \{s_{r+/-}\} \\ \geq 0}} \mathcal{C}(\Theta_{\text{LL}}) = \sum_{r^+ \in R^+} \Lambda_{r^+} s_{r^+} - \sum_{r^- \in R^-} \Lambda_{r^-} s_{r^-} \quad (4.1c)$$

$$\text{s.t. } \sum_{r^+ \in R^+} s_{r^+} - \sum_{r^- \in R^-} s_{r^-} = - (e^{\text{imb},+} - e^{\text{imb},-}) - \hat{\text{SI}} : \lambda^{\text{SI}} \quad (4.1\text{d})$$

$$s_{r^+} \leq S_{r^+} : \mu_{r^+} \quad \forall r^+ \in R^+ \quad (4.1\text{e})$$

$$s_{r^-} \leq S_{r^-} : \mu_{r^-} \quad \forall r^- \in R^- \quad (4.1\text{f})$$

where $\Theta_{\text{LL}} = \{s_{r^+}, s_{r^-}\}$ are the primal variables of the lower-level problem, and $\{\lambda^{\text{SI}}, \mu_{r^+}, \mu_{r^-}\}$ are the dual variables of the lower-level problem. The objective function (4.1a), maximizing the profit of the market actor, is computed as:

$$f^{\text{IS}}(e^{\text{imb},+/-}, \lambda^{\text{SI}}) = (\lambda^{\text{SI}} - C^+)e^{\text{imb},+} - (\lambda^{\text{SI}} - C^-)e^{\text{imb},-} \quad (4.2)$$

where C^+ and C^- define the cost structure of the asset. Practically, the market actor is incentivized to adopt a surplus energy position $e^{\text{imb},+}$ when $\lambda^{\text{SI}} > C^+$, and to favor a shortage energy position $e^{\text{imb},-}$ when $\lambda^{\text{SI}} < C^-$. Constraint (4.1b) ensures that these decisions comply with the technical margins Π_{UL} of the agent.

The lower-level reflects the costs minimization problem (4.1c), in which the TSO carries out the merit-order-based activation of the balancing energy products (where the more economic balancing energy bids are activated first). The offers $r^+ \in R^+$ and $r^- \in R^-$ are respectively activated at a price Λ_{r^+} and Λ_{r^-} to compensate the negative and positive imbalances. The imbalance price λ^{SI} is the price associated with the last (marginally activated) offer, which is endogenously obtained from the dual variable of the constraint (4.1d). The latter guarantees that the activated amount of reserves exactly offsets the anticipated imbalances caused by all other actors ($\hat{\text{SI}}$), while accounting for the strategic participation of the BRP ($e^{\text{imb},+/-}$). Finally, the set of constraints (4.1f) ensures that the activated balancing volumes s_{r^+} and s_{r^-} do not violate the energy limits (i.e., capacity S_{r^+} and S_{r^-} offered at an earlier stage). Note that the dual variables λ^{SI} , μ_{r^+} and μ_{r^-} represent the sensitivity of the lower-level objective function if the right-hand sides of their associated constraint are increased marginally. Practically, for instance, the value of the dual variable μ_{r^+} gives the amount by which having a marginal increase of the reserve capacity S_{r^+} would decrease the lower-level objective function's optimal value (i.e., the total cost). It can be interpreted as a shadow price of constraint (4.1e) for which the TSO would pay to have one unit more of the resource r^+ available.

4.3. From the Bi-level Model to a Mixed-Integer Linear Formulation

In order to solve the resulting bi-level problem (4.1) by off-the-shelf optimization solvers, it can be converted into an equivalent mixed-integer linear programming (MILP) formulation using the following steps [143]: i) the convex (linear and continuous) lower-level problem is replaced by its Karush-Kuhn-Tucker (KKT) optimality conditions, ii) the non-linearities of the complementarity conditions within the KKT conditions are equivalently expressed as a set of mixed-integer linear constraints using a Big-M approach, and iii) the bilinear term in the objective function is replaced by the related equivalent linear expression from the strong duality equation of the lower-level problem.

Karush-Kuhn-Tucker Optimality Conditions

The Karush-Kuhn-Tucker (KKT) conditions are conditions that the optimal solution of an optimization problem should satisfy. When the optimization problem is a linear program, the KKT conditions are necessary and sufficient for optimality [144]. Replacing the lower-level problem by its KKT conditions transforms the nested optimization problem (4.1) into a mathematical program with equilibrium constraints (MPEC), allowing to obtain a single level problem. The KKT conditions of the lower-level problem (4.1c)-(4.1f) are expressed as:

$$\sum_{r^+ \in R^+} s_{r^+} - \sum_{r^- \in R^-} s_{r^-} = -(e^{imb,+} - e^{imb,-}) - \hat{S}\mathbf{I} \quad (4.3a)$$

$$0 \leq -s_{r^+} + S_{r^+} \perp \mu_{r^+} \geq 0 \quad \forall r^+ \in R^+ \quad (4.3b)$$

$$0 \leq -s_{r^-} + S_{r^-} \perp \mu_{r^-} \geq 0 \quad \forall r^- \in R^- \quad (4.3c)$$

$$0 \leq s_{r^+} \perp \Lambda_{r^+} - \lambda^{SI} + \mu_{r^+} \geq 0 \quad \forall r^+ \in R^+ \quad (4.3d)$$

$$0 \leq s_{r^-} \perp -\Lambda_{r^-} + \lambda^{SI} + \mu_{r^-} \geq 0 \quad \forall r^- \in R^- \quad (4.3e)$$

where the \perp operator enforces the perpendicularity between the terms on the left-hand and right-hand sides, i.e., their product is zero.

Constraint (4.3a) and the left-hand sides of constraints (4.3b)-(4.3e) ensure the feasibility of the primal problem. The right-hand sides of constraints (4.3b)-(4.3c) enforce the feasibility of the dual problem, while the right-hand sides of constraints (4.3d)-(4.3e) are retrieved from the stationary conditions. Overall, the non-linear constraints (4.3b)-(4.3e) are the complementarity slackness conditions, written in a concise fashion.

Linearisation of the Complementarity Slackness Conditions

The big-M method [145] is used to transform the non-linear complementarity slackness conditions (4.3b)-(4.3e) into mixed-integer linear inequalities. It introduces, for each complementarity slackness condition, a large positive constant M and a binary variable z . The resulting MPEC model in Section 4.3 is thus recasted as a mixed-integer linear problem (MILP) as follows:

$$0 \leq -s_{r+} + S_{r+} \leq z_{r+}^1 M_{r+}^1 \quad \forall r^+ \in R^+ \quad (4.4a)$$

$$0 \leq \mu_{r+} \leq (1 - z_{r+}^1) M_{r+}^2 \quad \forall r^+ \in R^+ \quad (4.4b)$$

$$0 \leq -s_{r-} + S_{r-} \leq z_{r-}^1 M_{r-}^1 \quad \forall r^- \in R^- \quad (4.4c)$$

$$0 \leq \mu_{r-} \leq (1 - z_{r-}^1) M_{r-}^2 \quad \forall r^- \in R^- \quad (4.4d)$$

$$0 \leq s_{r+} \leq z_{r+}^2 M_{r+}^3 \quad \forall r^+ \in R^+ \quad (4.4e)$$

$$0 \leq \Lambda_{r+} - \lambda^{SI} + \mu_{r+} \leq (1 - z_{r+}^2) M_{r+}^4 \quad \forall r^+ \in R^+ \quad (4.4f)$$

$$0 \leq s_{r-} \leq z_{r-}^2 M_{r-}^3 \quad \forall r^- \in R^- \quad (4.4g)$$

$$0 \leq -\Lambda_{r-} + \lambda^{SI} + \mu_{r-} \leq (1 - z_{r-}^2) M_{r-}^4 \quad \forall r^- \in R^- \quad (4.4h)$$

The selection of appropriate values for the big-M parameters in the optimization problem can be a challenging task. When big-M values are too big, the resulting formulation may not hold the complementarity condition. On the other hand, too small big-M values may result in numerical ill-conditioning and/or increase drastically the computational time [92]. One way for choosing appropriate big-M values is first arbitrarily setting (for each of the complementarity conditions linearized) a large value for M , e.g., 10^7 , and then solve the model. In ex-post, the results are checked to investigate whether each of the complementarity conditions holds. If not, the value of the corresponding M is reduced until all complementarity conditions are satisfied [92]. This trial-and-error approach is specifically hard to apply in the context of our optimization problem, which is solved iteratively multiple times with varying parameters. Consequently, we have preferred to select adequate big-M values based on the economic or physical upper bounds of their associated variables [146]. The big-M values used in our case studies are shown in Table 4.1. For instance, the constraints related to the dispatch of the balancing energy products $s_{r+/-}$, i.e., eq. (4.4a), (4.4c), (4.4e), (4.4g), are naturally bounded by the energy limits of blocks $S_{r+/-}$. Consequently, the associated big-Ms values M_{r+}^1 and M_{r+}^3 (or M_{r-}^1 and M_{r-}^3) are fixed by the energy limits of block S_{r+} (or S_{r-}). The same logic can be applied for determining M_{r+}^2 , M_{r+}^4 , M_{r-}^2 and M_{r-}^4 .

M_j^i	$i = \{1, 3\}$	$i = 2$	$i = 4$
$j = r_+$	S_{r+}	$\max_{r^+ \in R^+} (\Lambda_{r^+}) - \Lambda_{r^+}$	$-\min_{r^- \in R^-} (\Lambda_{r^-}) + \Lambda_{r^+}$
$j = r_-$	S_{r-}	$-\min_{r^- \in R^-} (\Lambda_{r^-}) + \Lambda_{r-}$	$\max_{r^+ \in R^+} (\Lambda_{r^+}) - \Lambda_{r-}$

Table 4.1.: Big-Ms values

Linearisation of the Upper Level Objective Function

The bilinear term $\lambda^{\text{SI}}(e^{\text{imb},+} - e^{\text{imb},-})$ inside the upper level objective function can be linearised by applying the strong duality theorem to the lower-level problem's objective function [74]. When the optimization problem is convex, the strong duality theorem states that the objective functions of the primal and dual problems have the same value at the optimum. From this equality, an equivalent linear expression can be retrieved for the bilinear term.

$$\begin{aligned}
 \lambda^{\text{SI}}(e^{\text{imb},+} - e^{\text{imb},-}) &= -\lambda^{\text{SI}}\hat{\text{SI}} \\
 &\quad - \sum_{r^+ \in R^+} (S_{r^+}\mu_{r^+} + \Lambda_{r^+}s_{r^+}) \\
 &\quad + \sum_{r^- \in R^-} (-S_{r^-}\mu_{r^-} + \Lambda_{r^-}s_{r^-})
 \end{aligned} \tag{4.5}$$

The Deterministic Mixed-Integer Linear Formulation

The final deterministic optimization model is concisely written as follows:

$$\max_{\Theta^D} \text{Eq. (4.1a) (in a linearized format via (4.5))} \tag{4.6a}$$

$$\text{s.t. Eq. (4.1b)} \tag{4.6b}$$

$$\text{Eq. (4.3a)} \tag{4.6c}$$

$$\text{Eq. (4.4a) - (4.4h)} \tag{4.6d}$$

The model (4.6) optimizes the set of variables Θ^D , which contains the upper variables $\{e^{\text{imb},+}, e^{\text{imb},-}\}$, the lower primal $\{s_{r^+}, s_{r^-}\}$ and dual $\{\lambda^{\text{SI}}, \mu_{r^+}, \mu_{r^-}\}$ variables of the lower-level problem, as well as the binary variables $\{z_{r^+}^1, z_{r^-}^1, z_{r^+}^2, z_{r^-}^2\}$. However, relying on a deterministic decision-making tool for procuring real-time balancing services can be very risky. Indeed, in case of an erroneous estimation of the future system imbalance state, the BRP can adopt a position that aggravates the imbalance of the power system, thereby suffering important financial penalties. In the following Section, two methods for characterizing the uncertainty of the system imbalance are proposed: i) a




Figure 4.4.: Illustration of the uncertainty set and scenarios provided by the estimated empirical cumulative distribution function $\text{ECDF}(\cdot)$.

deterministic and set-based method, which is used for the robust optimization framework, and ii) a scenario-based method, which is applied in the stochastic programming framework.

4.4. Uncertainty Characterization

The outputs of the probabilistic forecasters in Chapter 3, i.e., the q -quantiles $\hat{\text{SI}}^{(q)}$ for $q \in \mathcal{Q} = \{0.05, 0.15, 0.25, 0.35, 0.45, 0.5, 0.65, 0.75, 0.85, 0.95\}$, such that $\mathbb{P}(\text{SI} \leq \hat{\text{SI}}^{(q)}) = q$, are leveraged for modeling the uncertainty of the system imbalance. First, at each quarter hour, the resulting discrete set of q -quantiles is used to estimate an empirical cumulative distribution function (ECDF) through cubic spline interpolation, which allows a continuous representation of the system imbalance distribution [147]. Then, based on the ECDF, the uncertainty is characterized depending on the subsequent stochastic decision-support tool. This procedure is illustrated in Fig. 4.4.

For the robust optimization framework, the uncertainty is characterized by an uncertainty set \mathcal{U} , which is defined as a box bounded by a symmetric pair of q -quantiles from the empirical ECDF. Hence, symmetric lower and upper bounds $\{\hat{\text{SI}}^{(q)}, \hat{\text{SI}}^{(1-q)}\}$ can be selected, each combination ensuring a certain probabilistic guarantee that the future system imbalance is realized within the

uncertainty set:

$$\mathcal{U} = \left\{ \text{SI} \in \mathbb{R}^1 : \hat{\text{SI}}^{(q)} \leq \text{SI} \leq \hat{\text{SI}}^{(1-q)} \right\}, \quad (4.7)$$

In this approach, a risk-aversion parameter $\epsilon = 1 - 2q$ can be defined for characterizing the size of the prediction interval associated with the uncertainty set. Hence, a larger ϵ yields a larger (and more conservative) uncertainty set, while a smaller value leads to a smaller uncertainty set. This risk-aversion parameter allows the BRP to adjust its risk policy, i.e., the degree of conservativeness of its decisions, by varying the size of the uncertainty set [148]. The economic interest of different sizes of uncertainty set is touched in Section 4.7, and more thoroughly investigated in Section 4.9.

For stochastic programming framework, the methodology presented in [147] is followed. A set of system imbalance scenarios $\{\hat{\text{SI}}_\omega, \forall \omega \in \{1, \dots, N\}\}$ is generated by applying the inverse transform method on the estimated ECDF. This consists in firstly drawing a random variable between $[0.05, 0.95]$ (the orange diamonds in Fig. 4.4) for subsequently retrieving a system imbalance scenario $\hat{\text{SI}}_\omega$ (the blue diamonds in Fig. 4.4). This Monte Carlo sampling approach is performed at each optimization period for a set Ω of $N = 100$ scenarios, each of them associated with a probability of occurrence $p = 1/N$. In our experiments, increasing further the number of scenarios does not improve the results of the stochastic programming optimization decision-support tool. If each objective outcome associated with each $\hat{\text{SI}}_\omega$ is effectively weighted given the predefined probabilities $p = 1/N$, the scenario-based optimization procedure is called risk-neutral. For risk-aware approach, the probabilities of each scenario outcome can be adjusted according to the BRP risk preferences through a risk measure (note that the sum of the risk-adjusted probabilities is still equals to 1). In this Chapter, this option is investigated via the Conditional Value-at-Risk (CVaR) metric (see Section 4.6), which puts more weight on the riskier scenario outcomes (that deviates adversely from the expectation) in the scenario-based optimization procedure.

4.5. The Robust Optimization Model

Robust optimization aims at computing a solution that is feasible for any realization within the uncertainty set \mathcal{U} , and that is optimal against the worst-case realization. The robust counterpart of the model (4.6) illustrated in Fig. 4.5 is written as:

$$\max_{\Theta^R} \min_{\text{SI} \in \mathcal{U}} f^{\text{IS}}(e^{imb, +/-}, \lambda^{\text{SI}}) \quad (4.8a)$$

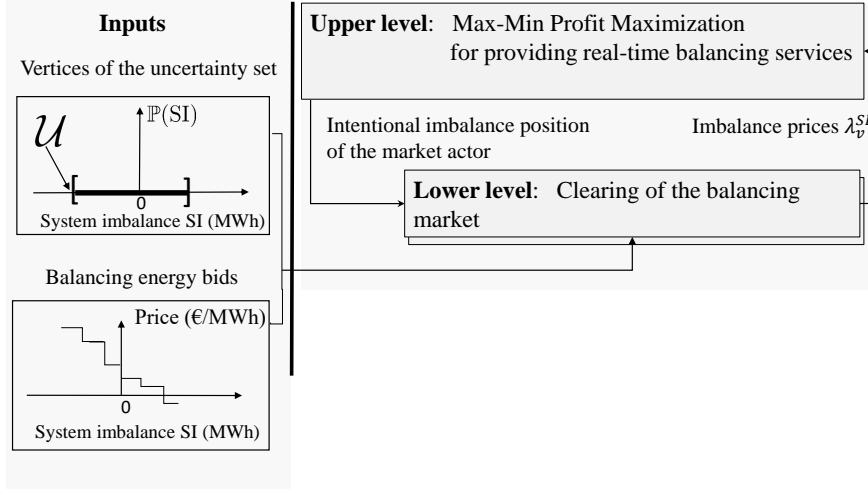


Figure 4.5.: The robust optimization framework.

$$\text{s.t. } e^{\text{imb},+/-} \in \Pi_{\text{UL}} \quad (4.8\text{b})$$

$$\text{Eq. (4.3a), } \forall \hat{\text{SI}} \in \mathcal{U} \quad (4.8\text{c})$$

$$\text{Eq. (4.4a) – (4.4h), } \forall \hat{\text{SI}} \in \mathcal{U} \quad (4.8\text{d})$$

where Θ^R is the set of variables $\{e^{\text{imb},+}, e^{\text{imb},-}, s_{r+}, s_{r-}, \lambda^{\text{SI}}, \mu_{r+}, \mu_{r-}, z_{r+}^1, z_{r-}^1, z_{r+}^2, z_{r-}^2\}$.

In this robust approach, the objective function (4.8a) immunizes the operational strategy of the BRP against the worst-case realization of the system imbalance contained in the uncertainty set \mathcal{U} . For a polyhedral uncertainty set (as in our case), it has been shown that the worst-case realization of the uncertainty set is located at one of its vertices [149]. This property allows to reformulate the robust problem (4.8) to be solved by off-the-shelf optimizers. Practically, an auxiliary variable σ is added for representing the worst-case profit $f^{\text{IS}}(\cdot)$ through the additional constraint (4.9b), and the continuity of the uncertainty set \mathcal{U} is managed by enumerating the finite number of vertices $v = \{V_1, V_2\}$ contained in \mathcal{U} . The robust optimization-based equivalent of problem reads as:

$$\max_{\sigma, \Theta^{RF}} \sigma \quad (4.9\text{a})$$

$$\text{s.t. } \sigma \leq f^{\text{IS}}(e^{\text{imb},+/-}, \lambda_v^{\text{SI}}), \quad \forall v = \{V_1, V_2\} \in \mathcal{U}, \quad (4.9\text{b})$$

$$\text{s.t. } e^{\text{imb},+/-} \in \Pi_{\text{UL}} \quad (4.9\text{c})$$

$$\text{Eq. (4.3a), } \forall v = \{V_1, V_2\} \in \mathcal{U}, \quad (4.9\text{d})$$

$$\text{Eq. (4.4a) - (4.4h)}, \quad \forall v = \{V_1, V_2\} \in \mathcal{U}, \quad (4.9e)$$

where Θ^{RF} is the set of variables $\{e^{\text{imb},+}, e^{\text{imb},-}, s_{v,r+}, s_{v,r-}, \lambda_v^{\text{SI}}, \mu_{v,r+}, \mu_{v,r-}, z_{v,r+}^1, z_{v,r-}^1, z_{v,r+}^2, z_{v,r-}^2\}$.

It should be noted that solving the worst-case realization is facilitated by the fact that the uncertainty set $\mathcal{U} \in \mathbb{R}^1$ in our robust formulation. Indeed, the worst-case realization is generally the lower predicted q -quantile of the system imbalance, except when the vertices of the uncertainty set are above and below zero. In the latter case, the robust approach will just prevent the provision of real-time balancing services. Hence, the determination of the worst-case realization can be externalized by a if-then procedure, for which the optimal real-time balancing service can be provided through the deterministic model (4.6).

4.6. The Stochastic Programming Optimization Model

Scenario-based stochastic programming is a well-known technique that optimizes the expected value of the objective function by representing the distribution of the uncertainty through a set of scenarios, each assigned with a probability. The scenario-based stochastic program (illustrated in Fig. 4.6) is written in model (4.10) as follows:

$$\max_{\Theta^{SP}} \mathbb{E} [f^{\text{IS}}(e^{\text{imb},+/-}, \lambda_w^{\text{SI}})] = \frac{1}{N} \sum_{w \in \Omega} f^{\text{IS}}(e^{\text{imb},+/-}, \lambda_w^{\text{SI}}) \quad (4.10a)$$

$$\text{s.t. } e^{\text{imb},+/-} \in \Pi_{\text{UL}} \quad (4.10b)$$

$$\text{Eq. (4.3a)}, \quad \forall w \in \Omega, \quad (4.10c)$$

$$\text{Eq. (4.4a) - (4.4h)}, \quad \forall w \in \Omega, \quad (4.10d)$$

where Θ^{SP} is the set of variables $\{e^{\text{imb},+}, e^{\text{imb},-}, s_{w,r+}, s_{w,r-}, \lambda_w^{\text{SI}}, \mu_{w,r+}, \mu_{w,r-}, z_{w,r+}^1, z_{w,r-}^1, z_{w,r+}^2, z_{w,r-}^2\}$.

The objective function (4.10a) consists of the expected profit of the market actor among the system imbalance scenarios $\{\hat{\text{SI}}_w, \forall w \in \Omega\}$. The balancing clearing process, performed for each scenario ω , is enforced by Eq. (4.10c)-(4.10d). When the decision-making process is repetitive, stochastic programming allows to better optimize the expected profit in a long run, but only if the probability of the distribution is well captured. To do so, the discrete set of scenarios has to be as much as representative of the original

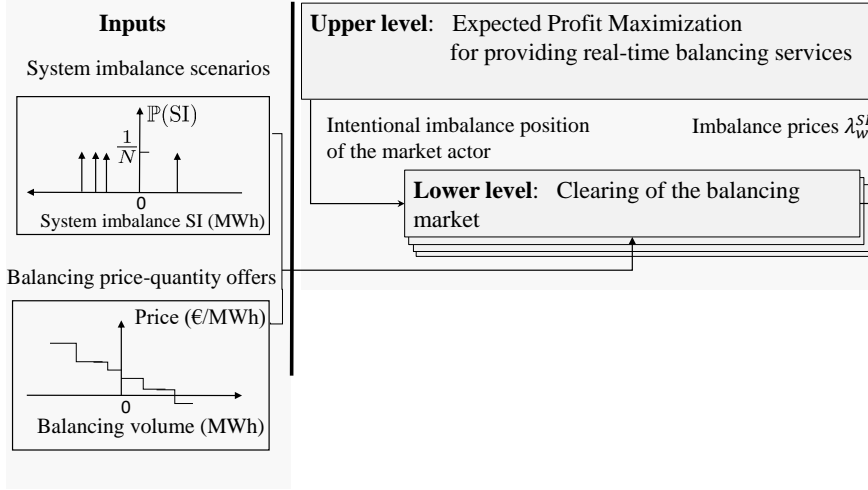


Figure 4.6.: The scenario-based stochastic programming optimization framework.

(continuous) uncertain distribution. However, one should notice that the constraints (4.10c)-(4.10d) are imposed for each scenario, increasing rapidly the dimension of the problem. Hence, when using stochastic programming, there is always a trade-off between the number of scenarios, allowing an accurate representation of the uncertainty, and the computational cost of the problem.

This gap between the original (continuous) uncertain distribution and the discrete set of scenarios can provide ex-post disappointments to the decision maker when its actual (true) objective outcome is revealed. One way for addressing this issue is to nuance the decisions of the risk-neutral model (4.10) by incorporating a notion of financial risk through the Conditional Value-at-Risk measure. This allows quantifying a level of trading risk, and to reduce the volatility of the profit among the set of scenarios Ω . Let $\epsilon \in [0, 1]$ representing a risk-aversion parameter, the CVaR_ϵ is defined as the expected profit of the $(1 - \epsilon) \times 100\%$ worst scenarios. The CVaR-based stochastic program can be defined as:

$$\max_{\substack{\zeta, \eta_w, \Theta \\ \geq 0}} \zeta - \frac{1}{1 - \beta} \sum_{w \in \Omega} \pi_w \eta_w \quad (4.11a)$$

$$\text{s.t. } \zeta - f^{\text{IS}} \left(e^{\text{imb}, +/ -}, \lambda_w^{\text{SI}} \right) \leq \eta_w, \quad \forall w \in \Omega, \quad (4.11b)$$

$$e^{\text{imb},+/-} \in \Pi_{\text{UL}} \quad (4.11\text{c})$$

$$\text{Eq. (4.3a)}, \quad \forall w \in \Omega, \quad (4.11\text{d})$$

$$\text{Eq. (4.4a) – (4.4h)}, \quad \forall w \in \Omega, \quad (4.11\text{e})$$

where ζ is the Value-at-Risk (VaR), i.e., the $(1-\epsilon)$ -quantile of the profit distribution of scenarios. The non-negative auxiliary variable η_w corresponds to the difference between the VaR ζ and the market actor profit $f^{\text{IS}}(\cdot)$ if it is positive.

In this risk-aware formulation, a larger risk-aversion parameter ϵ leads to more conservative decisions (as only the few worst scenarios are considered), while a smaller value entails more risk-neutral decisions (as the set of scenarios is widened). Note that a trade-off between the expected profit and the CVaR metric, i.e., an adjusted-risk approach weighting both contributions in the objective function, is also doable. The economic interest of risk awareness in stochastic programming for our application is investigated in Section 4.9.

4.7. Economic Interest of Accurate Probabilistic Forecasts

The economic interest of generating more accurate probabilistic forecasts is investigated in this Section by studying their economic gains when associated with the robust model (4.9). Practically, the lower-level problem of the robust model (4.9) is cleared based on symmetric pairs of the predicted quantiles (which constitute the uncertainty set). Hence, different pairs of quantiles are investigated for different prediction models on the basis of the following uncertainty sets $\mathcal{U} \in \{[\hat{\text{SI}}^{(0.05)}, \hat{\text{SI}}^{(0.95)}], [\hat{\text{SI}}^{(0.15)}, \hat{\text{SI}}^{(0.85)}], [\hat{\text{SI}}^{(0.25)}, \hat{\text{SI}}^{(0.75)}], [\hat{\text{SI}}^{(0.35)}, \hat{\text{SI}}^{(0.65)}], [\hat{\text{SI}}^{(0.45)}, \hat{\text{SI}}^{(0.55)}]\}$. This allows to quantify the economic gain that could be obtained by more accurate q -quantile SI predictions on real-life market data. The ex-post economic profits of the BRP $f^{\text{OS}}(\cdot)$ are computed via an ex-post out-of-sample analysis, whose imbalance price is obtained via the stand-alone clearing of the balancing energy market (4.1c)-(4.1f) based on the actual realization of the SI and the imbalance position $e^{\text{imb},+/-}$ of the BRP (obtained via the robust-based optimization). The overall procedure is run sequentially at the start of each quarter-of-an-hour of the test set, i.e., January and February 2018. In accordance with the methodology previously presented in Chapter 3, the following probabilistic forecasting models are assessed:

- the quantile regression forest (QRF);

- the gradient boosting regression tree (QGBRT);
- the 1-layer feed-forward neural network (FFNN);
- the stacked feed-forward neural network (S-FFNN);
- the Long Short Term Memory recurrent neural network (LSTM);
- the Bidirectional Long Short Term Memory recurrent neural network (BLSTM);
- the sequence-to-sequence model (Seq2Seq).

Note that only the time periods of the dataset are different from Chapter 3. In this Chapter, the forecasters are trained using historical data from 2014 until end of December 2016 for outputting the q -quantiles $\{\text{SI}^{(q)}, \forall q = \{0.05, 0.15, 0.25, 0.35, 0.45, 0.5, 0.55, 0.65, 0.75, 0.85, 0.95\}\}$. The year 2017 is used as a validation set to select the hyper-parameters of the different models. Then, one step-ahead probabilistic prediction of the system imbalance is inferred at the start of each quarter hour of the test set (i.e., January and February 2018), for which their economic interest is assessed.

The probabilistic performances over the entire test are presented in Table 4.2 for the different models. Two scoring metrics are reported: i) the pinball loss, which is averaged across all q -quantiles of interest, and ii) the Winkler score, which is calculated for the prediction intervals $\beta = \{0.1, 0.3, 0.5, 0.7, 0.9\}$. The best results are highlighted in bold font. Similar observations with respect to Chapter 3 can be reported. The Seq2seq model, whose architecture is designed to better exploit both past observed and future known dynamic information, outperforms all other machine learning models. The QRF model seems to achieve lower performances than other models, while the QGBRT models (one for each q -quantile) provide similar, even better, results w.r.t. the FFNN architectures. Note that the S-FFNN (with 3 hidden layers) still obtains better results than the shallow FFNN.

The probabilistic predictions are then incorporated into the robust dispatch strategy, where the quality of the resulting decisions is compared in Table 4.3. As a measure of the upper bound of the profit that can be generated, the decision-support model was firstly run with the perfect knowledge of the system imbalance. Then, the aggregated profits (summed over all quarter hours of the 2-months test set period) $\mathbb{E}(f^{\text{IS}})$ that were expected at the end of the optimization procedure are summarized in the upper part of the Table 4.3. These (in-sample) results are then put into perspective with the actual profits $\mathbb{E}(f^{\text{OS}})$ given in the lower part of Table 4.3. These results are obtained via the

Table 4.2.: Pinball loss and Winkler scores of the forecasting methods on January and February 2018 for the one-step ahead probabilistic prediction of the system imbalance. The Winkler score is expressed for the prediction intervals $\beta = \{0.1, 0.3, 0.5, 0.7, 0.9\}$.

Topology	Pinball loss [MW]	Winkler score [MW]				
		0.1	0.3	0.5	0.7	0.9
QRF	223	332	246	200	166	137
FFNN	207	310	227	185	154	127
S-FFNN	199	299	220	178	148	123
QGBRT	189	309	212	168	138	114
Seq2seq	179	268	197	159	132	109

(ex-post) clearing of the imbalance settlement, based on the optimized position $e^{\text{imb},+}$ and $e^{\text{imb},-}$ of the BRP and the actual realization of the system imbalance. For focusing on the accuracy impact of the probabilistic forecasts, the feasible region of the BRP portfolio Π_{UL} in the robust model is only constrained by (upward and downward) 120 MW power limits at each quarter hour. The operating costs $\{C^+, C^-\}$ are respectively set at 50 and 30 €/MWh, which ensures that the BRP will exploit consistent price spreads. Practically, the BRP will adopt a surplus energy position, i.e., $e^{\text{imb},+} > 0$, only if $\lambda^{\text{SI}} > 50$ €/MWh, and a shortage position, i.e., $e^{\text{imb},-} > 0$, only if $\lambda^{\text{SI}} < 30$ €/MWh. We consider that no contracts were agreed in advance, and that only the decisions $e^{\text{imb},+/-}$ describe the imbalance position of the BRP. The resulting MILP model is implemented in Python, using the Pulp library, and solved using the Gurobi 8.1.1 solver. It is important to mention that the computational time is always lower than 1 minute (whereas the forecasting tool necessitates less than 1 second to provide the probabilistic forecasts of interest), which suggests that the robust formulation can be exploited in the real-time operation of a BRP with flexible resources.

Based on Table 4.3, several trends can be identified. Firstly, it is observed that the quality of predictions constitutes a prominent factor to take reliable decisions. In this way, using the best prediction interval for each technique (denoted in bold in the lower part of Table 4.2), the Seq2seq model increases the operational profit between 40 k€ (compared to the second best model, i.e., QGBRT) and 298 k€ (for QRF), which corresponds to relative increases of respectively 2.8% and 21% throughout the test period of January-February 2018. In this way, even small improvements in the prediction accuracy can result into significant additional profits. This strongly paves the way to further research to enhance prediction tools. In this way, the best predictor yields

Table 4.3.: Operational profits of different forecasters describing different bounds of the uncertainty set when providing robust real-time balancing services.

	Aggregated expected (ex-ante) profits over the test set [k€]					
Perfect forecast	4201					
	Bounds 05-95	Bounds 15-85	Bounds 25-75	Bounds 35-65	Bounds 45-55	Bound 50
QRF	161	554	1047	1624	2316	2722
FFNN	359	882	1396	1859	2533	2975
S-FFNN	587	1222	1704	2258	2784	3136
QGBRT	216	954	1574	2111	2697	3070
Seq2seq	574	1213	1765	2315	2959	3315
	Aggregated actual (ex-post) profits over the test set [k€]					
Perfect forecast	4201					
	Bounds 05-95	Bounds 15-85	Bounds 25-75	Bounds 35-65	Bounds 45-55	Bound 50
QRF	480 939	1131	1019	634	245	
FFNN	850	1126	1203	1015	637	264
S-FFNN	1040	1282	1184	971	564	332
QGBRT	798	1247	1389	1228	950	663
Seq2seq	1151	1429	1391	1251	1023	847

an operational ex-post profit of 1429 k€, whereas perfect forecasts would generate 4201 k€. Moreover, in accordance with its forecast outcomes (Table 4.2), QGBRT is very competitive for tightened quantiles, but leads to a lower performance for larger quantiles (since the Winkler score is typically high for the prediction interval $\beta = 0.1$).

Secondly, the robust formulation ensures that the portfolio never participates in the single price imbalance settlement when the bounds of the uncertainty set (i.e., the pair of the q -quantiles $\hat{SI}^{(q)}$) are of different signs. On the contrary, when the predicted bounds are sign-consistent, the market player may deviate from its balanced position with the objective to help at restoring the power system balance. For this to happen, the imbalance price must be sufficiently attractive to cover all intrinsic costs of the unit (which is true in most cases due to the extreme price regimes of the imbalance price). In this way, in the deterministic setting (where only the predicted median $\hat{SI}^{(0.5)}$ is considered), the market player almost systematically participates in the imbalance settlement, which represents an amount of 75 255 MWh over the test set period (when the Seq2seq is used as forecaster). By contrast, the uncertainty set $\mathcal{U} = [\hat{SI}^{(0.05)}, \hat{SI}^{(0.95)}]$ leads to a moderate contribution of 16 352 MWh (which are played 27% of the time, i.e., an average of 25 times a day for the 96 daily settlement periods).

Thirdly, we see that the expected and actual profits can be very different. Such discrepancies are exacerbated for the tightened quantiles since the decisions are then based on forecasts with low reliability. Such aggressive approaches tend to overestimate the profit that will be actually generated, leading to ex-post disappointments when the actual outcome is revealed. For instance, the approach with the 45-55 quantiles represents a reliability of 10% that the actual system imbalance lies in the prediction interval, which ultimately jeopardizes the performance of the optimization. This illustrates the importance of incorporating risk-awareness in the decision-support tool of the BRP to hedge against such situations.

In Table 4.3, the most suited approach is the optimization performed with the 15-85 quantiles (obtained with the Seq2seq), since this tool leads to the highest actual profits, and therefore to the best trade-off between conservativeness and economic performance. This strategy is sufficiently audacious to properly take advantage of favorable situations, while hedging against the inherent volatility of the system imbalance signal, which avoids to participate when the market conditions are unsure. Indeed, results highlight that the 15-85 quantiles approach leads to erroneous offers (that infer financial penalties)

occurring 6.5% of the time. In comparison, the deterministic (quantile 50) and aggressive strategies (45-55, 35-65 and 25-75 quantiles) result in respectively 38%, 32%, 22% and 14% of erroneous offers, while such misinformed decisions happen respectively 1% of the time for the more conservative strategies (05-95 quantiles).

4.8. Economic Interest of the Bi-level Structure

Two types of approach can be envisaged for modeling the participation of a market player in electricity markets: the price-taker and the price-maker approaches. A price-taker market player accepts an exogenous price signal, and participates in the electricity markets without affecting the price signal, while a price-maker market player may influence the market-clearing outcome by its own participation. In our electricity market application, i.e., the European single price imbalance settlement, there is a high sensitivity of the imbalance price with respect to the (very fluctuating) system imbalance. Hence, each actor (even small ones) may act as a price-maker in the single price imbalance settlement, and may incur sharp regime switching in the imbalance price (recalling Fig. 4.2). In this section, the economic interest of embedding the market-clearing process of the balancing energy market within the robust optimization model is challenged. To that end, the robust optimization model is solved in the same conditions (based on the predictions given by the Seq2seq) in a two-step approach. Practically, the lower-level (4.9) is firstly cleared independently (disregarding the actions of the market player) for the pair of q -quantiles $[\hat{SI}^{(q)}, \hat{SI}^{(1-q)}]$ constituting the uncertainty set, which provides two estimations of the future imbalance prices $[\hat{\lambda}_{V_1}^{SI}, \hat{\lambda}_{V_2}^{SI}]$. Then, the resulting imbalance prices are treated as parameters (and **not** decision variables) in the upper level optimization problem (4.9a)-(4.9b) (without considering the lower-level problem). The expected and actual profits for both approaches are shown in Table 4.4.

Results indicate that the price-taker assumption is not appropriate to model the close-to-real-time participation of BRPs in the single price imbalance settlement. Indeed, when the forecasted quantiles of the system imbalance are of the same sign, the actor provides its full capacity (regardless of its impact on the system imbalance), which often results in switching the system conditions. The market actor then consumes at high prices, and produces energy for low prices, which ultimately leads to negative profits (especially for aggressive approaches). In this way, the price-taker assumption is systematically over-optimistic, which tends to result in very inefficient strategies when the actual outcome (ex-post profit) is revealed. The price-maker assumption, on the other

Table 4.4.: Comparison of the profit generated by price-maker and price-taker assumptions for different risk-attitudes

	Aggregated expected profits		Aggregated ex-post profits	
	over the test set [k€]		over the test set [k€]	
	Price-taker	Price-maker	Price-taker	Price-maker
Bounds 05-95	1468	574	773	1151
Bounds 15-85	3142	1213	-394	1429
Bounds 25-75	4370	1765	-1588	1391
Bounds 35-65	5458	2315	-2870	1251
Bounds 45-55	6598	2959	-4396	1023
Bound 50	7296	3315	-5147	847

hand, allows to better hedge against the real-time volatility of the power system conditions, by properly considering the impact of the BRP decisions.

4.9. Economic Interest of Risk Awareness for Providing the Real-Time Balancing Services

In this Section, the economic interest of including a financial risk management in the stochastic decision-support tool of the BRP is studied. Following the robust model (4.8), the financial risk can be adjusted by varying the size of the uncertainty set. Note that Table 4.3 has already performed a similar analysis, but this section goes further by extending the test set over the entire year 2018. On the other hand, the financial risk in the scenario-based stochastic programming model (4.11) is managed by including the conditional value-at-risk (CVaR) in the objective function. This risk approach assigns higher probabilities to the scenarios with lowest profits. The economic results are computed over the year 2018 based on the probabilistic forecasts of the Seq2seq model (which are also extended). The notation RO stands for the robust optimization model, while SP-CVaR denotes the CVaR-based stochastic optimization model.

The economic results in Table 4.5 of both risk-aware optimization methods are differentiated according to their respective risk-aversion parameter ϵ . More specifically, we consider 11 risk policies for the BRP corresponding to the following $\epsilon \in \{0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 0.98\}$ risk-aversion

Table 4.5.: Average of the out-of-sample profits $f_\epsilon^{OS}(\cdot)$ for the different risk policies over the 1-year period.

Model	Unit [€/ Δt]	Risk-aversion parameter (ϵ)										
		0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.98
RO	$\mathbb{E}(f_\epsilon^{OS})$	124.1	185.1	227.4	251.8	263.4	270.6	275.4	281.2	280.5	221.2	100.5
SP-CVaR	$\mathbb{E}(f_\epsilon^{OS})$	338	335	331,9	327	315	300,1	292,1	283,2	244,2	175,7	129

parameters. This allows covering a wide range of possible risk policies, from the most risk-seeking ($\epsilon = 0$) to the most risk-averse ($\epsilon = 0.98$) ones. As a reminder, the risk-aversion parameter ϵ in the robust model is defined as $\epsilon = 1 - 2q$, which characterizes the size of the prediction interval associated with the uncertainty set. Besides, the risk-aversion parameter ϵ in the CVaR metric defines the $(1 - \epsilon) \times 100\%$ of worst-case scenarios on which the scenario-based stochastic program is optimized. Table 4.5 presents only the out-of-sample performance of the two optimization models for all possible risk policies ($E = 11$) over the entire test set.

For RO, the results of Table 4.5 confirms that the risk policy $\epsilon = 0.7$, which is described by an uncertainty set $\mathcal{U} = [\hat{SI}^{0.15}, \hat{SI}^{0.85}]$, is the most adequate risk profile for the BRP. It is also worth mentioning that the SP-CVaR approach yields the best performance for the risk-neutral strategy, which mainly stems from two reasons: i) the Seq2seq forecaster provides ‘high-quality’ scenarios of the system imbalance, and ii) this framework optimizes the profit in expectation, which is efficient for decision-making procedure that occur very regularly in time. In particular, the risk-neutral SP-CVaR outperforms the robust optimization method, with a relative increase of the actual profits of around 20% on average at their optimal risk-aversion parameter. The ex-post economic profits for the SP-CVaR then decrease gradually with the rise of conservativeness of the CVaR risk-aversion parameter.

4.10. Conclusion

In this Chapter, we present the market opportunities and risks associated with the provision of real-time balancing services in the single price imbalance settlement. A decision-support tool relying on the bi-level methodology is firstly proposed for participating in such a mechanism, allowing to capture the interaction between the real-time balancing service of the BRP and the clearing of the balancing energy market. The resulting problem involves an optimization problem constrained by another optimization problem. Mathematical reformulation steps are thus exposed to convert the nested

bi-level model into a mixed-integer linear programming (MILP) problem. The final (deterministic) linearized model is used as a basis for developing risk-aware stochastic decision-support tools.

Indeed, two different risk-aware stochastic decision-support tools are provided: i) a robust optimization model, and ii) a scenario-based stochastic programming, where the notion of financial risk management is included via the conditional Value-at-Risk measure (CVaR). Both risk-aware decision support tools are assessed and compared in extensive case studies using real-life market data from the Belgian power systems.

Outcomes suggest that i) gains of accuracy in the probabilistic predictions allow achieving better decisions, and thus, better ex-post economic profits, ii) the bi-level structure is efficient for hedging against the inherent small volume of system imbalances and the associated imbalance price regime switching effect, iii) the determination of the uncertainty set is not straightforward for the robust approach, as the ex-post economic profits vary widely depending the size of the uncertainty set, and iv) the risk-aware stochastic program shows the best ex-post economic profits for the risk-neutral approach, which decreases gradually with the rise of conservativeness of the CVaR risk-aversion parameter.

Regarding risk awareness, the next chapter investigates whether an automatic risk-adjusted approach can be designed for dynamically changing the risk profile of the BRP between consecutive market periods.

Chapter Publication

- **J. Bottieau**, L. Hubert, Z. De Grève, F. Vallée and J-F. Toubéau, "Very-Short-Term Probabilistic Forecasting for a Risk-Aware Participation in the Single Price Imbalance Settlement, " in *IEEE Trans. Power Syst.*, vol. 35, no. 2, pp. 1218-1230, 2020.

CHAPTER 5.

Automatic Risk-Adjusted Provision of Real-Time Balancing Services

Trading strategies in short-term electricity markets generally employ risk awareness for reducing, *inter alia*, their exposure to the volatility of electricity prices. To ensure an optimal balance between risk and profit, risk-aversion parameters are traditionally fine-tuned via an offline out-of-sample analysis. Such a computationally-intensive analysis is typically run once (e.g., the one performed in Section 4.9), which yields time-invariant risk policies. Instead, this Section proposes the use of Machine Learning to select, in an online fashion, optimal risk-aversion parameters. This novel automatic risk-tuning approach offers the benefit of continuously adjusting the risk policy based on the dynamically changing market operating conditions. The proposed approach is tested on two risk-aversion parameters, i.e., the budget of uncertainty and the confidence level of the conditional value-at-risk, respectively considering the robust optimization and CVaR-based stochastic programming frameworks. Both automatic risk-adjusted decision-support tools are then assessed and tested on real-world market data from the Belgian power system.

This Chapter is organized as follows. Section 5.1 provides an overview about how financial risk is traditionally managed in electricity markets trading strategies. Section 5.2 describes i) the goal of the automatic risk adjustment tool, and how it fits in the risk-aware decision-support tools of Chapter 4, ii) how the automatic risk-adjustment is trained and deployed, and iii) the different machine learning models used. In Section 5.3, time-invariant and automatic risk policies are assessed and compared in the Belgian case study. Finally, conclusions are presented in Section 5.4.

5.1. Risk Awareness in Electricity markets Trading Strategies

While competing in liberalized electricity markets, actors adapt their short-term dispatch decisions based on their expectations of future market outcomes to maximize their profit [30]. These decisions are made while facing market uncertainties stemming from, e.g., the prevailing market prices, which expose the actor to financial risks [150]. In this context, stochastic decision support tools, including financial risk management, allow improved scheduling decisions in short-term electricity markets, giving the possibility for actors to manage the risks associated with their positions [69].

Two distinct methodologies can be considered for trading strategies: i) *performance satisfying methods*, such as information-gap theory decision [151], [152], which ensure a minimum acceptable profit, and ii) *performance maximization methods*, which, in contrast, maximize the expected profit of the market actor given its representation of the uncertainty space. Typical examples of performance maximization methods are robust optimization [148], [153] and stochastic programming [154]–[156], which were used in Chapter 4. In performance maximization methods, the market actors take two successive decisions when employing risk awareness: i) the selection of their risk attitude, i.e., the setting of their financial risk management (e.g., their own risk averseness level), which determines their risk policy, and ii) their scheduling/dispatch decision, which maximizes their expected profit given their risk policy.

Risk awareness in performance maximization methods has attracted a high-level of interest within the power systems community, including several studies on conventional electricity generation [157]–[159], energy storage systems [160], [161], weather-dependent generation [154], [155], [162], [163], demand-response [164]–[166], and hybrid power plants [148], [153], [156], [167]–[170]. Overall, robust optimisation-based approaches allow to adjust the conservativeness of the decisions by varying the budget of uncertainty, i.e., the size and shape of the uncertainty set [148]. On the other hand, following a scenario-based stochastic optimization framework, most authors incorporate risk measures, e.g., the conditional value-at-risk (CVaR), in the objective function to assign higher weights to the scenarios with lowest profits [171].

Currently, existing research efforts in performance maximization methods focus on evaluating the optimal risk policy through an offline process. This process entails varying the risk-aversion parameters – e.g., the confidence level

of the CVaR metric or the budget of uncertainty – on extensive out-of-sample evaluations to select the optimal risk attitude – see e.g., [165], [166], [170], [172]. Such a computationally-intensive analysis is typically run once, which does not allow capturing the dependency of the optimal risk policy on the dynamically changing market operating conditions. In this sense, in all the aforementioned approaches, there is – to the best of the author knowledge – no systematic way to autonomously and dynamically adjust the selection of the risk-aversion parameters. Indeed, the out-of-sample analysis typically yields only time-invariant risk policies, even though the real-time power system conditions, and, as a result, the financial risk and the optimal risk policy, may significantly vary over time.

This Chapter aims at addressing this limitation by leveraging Machine Learning (ML) techniques to adjust, at each decision step, the risk policy of an actor based on the current state of expected market outcomes. More specifically, an ML-based module is designed to estimate the time-specific out-of-sample economic performance of different risk attitudes *based on past trading sessions*, allowing to determine autonomously the most optimal online risk attitude. Practically, the proposed automatic risk-adjusted decision-support tool is applied on the provision of real-time balancing services via the single price imbalance settlement mechanism. The risk policy is crucial in this application since the system conditions are highly volatile and difficult to predict. In addition, actors are exposed to significant financial penalties in case of sub-optimal decisions. In this regard, the risk-aware stochastic formulations proposed in Chapter 4 are extended to an automatic risk-adjusted framework, which considers both CVaR-based stochastic and robust optimization frameworks.

5.2. Automatic Risk-Adjusted Decision-Support Framework

The electricity market application is the one presented in Chapter 4, where a Balance Responsible Party (BRP) has the ability to deviate from its energy position to support the real-time system balancing within a single imbalance pricing scheme. In Chapter 4, two risk-aware stochastic decision-support tools were developed, namely the robust optimization model (4.9) and the CVaR-based stochastic program (4.11), aiming at reducing the exposure of the BRP to the significant variability and uncertainty exhibited by the system imbalance.

Hence, based on a risk attitude ϵ , the risk-aware decision-support tool aims at providing an optimal imbalance position for the BRP, while indicating the associated in-sample objective outcome f_ϵ^{IS} for the next market period. In




Figure 5.1.: The integration of the proposed automatic risk adjustment tool within the decision-making process of an actor participating in the single price imbalance settlement. The automatic risk-adjusted decision-support tool is run sequentially (96 times a day) at the start of each 15 minutes imbalance settlement period.

this Section, as shown in Fig. 5.1, the decision-making process of the BRP is further supported by adding an automatic risk adjustment tool. This tool relies on a Machine Learning (ML)-based approach to adjust the risk attitude of the BRP. The selected risk attitude is then used in the subsequent risk-aware stochastic optimization tool to calculate its optimal position in the imbalance settlement market. Once the imbalances of all BRPs are settled, an out-of-sample objective outcome f_{ϵ}^{OS} can be generated in an ex-post analysis by confronting the optimised imbalance position of the market actor at a given risk attitude ϵ with the actual realization of the system imbalance. This computation can stem from real-life market outcomes or complex market simulators. The latter option is used in the case studies.

5.2.1. Goal of the Automatic Risk Adjustment Tool

There exists inevitably a gap between an in-sample objective outcome (f_{ϵ}^{IS}) provided by a model, i.e., the optimal objective value of a risk-aware stochastic optimization tool at a risk-aversion parameter ϵ , and the corresponding out-of-sample objective outcome (f_{ϵ}^{OS}) obtained when the actual realization of the uncertainty is revealed. The gap between in-sample and out-of-sample objective outcomes can arise from either a misrepresentation of the uncertainty or a simplified representation of the market environment. Indeed, a risk-aware stochastic optimization tool is generally composed of two stages: i) an uncertainty model, which allows identifying a probabilistic representation of the future possible realizations of random variables (see Section 4.4 of Chapter 4), and ii) a risk-aware stochastic optimization model, which mathematically expresses the market environment in which the actor operates, and generates optimal decision outcomes at a pre-defined risk attitude, based on the prior

representation of uncertainties (see Sections 4.5-4.6 of Chapter 4). In our market application, the misrepresentation of the uncertainty may come from i) a bad probabilistic forecaster, or ii) the set-based or discrete characterization of the uncertainty pertaining to the robust and scenario-based optimization approaches, respectively. Besides, the market environment assumes that: i) the imbalance price is constant over intervals of 100 MW, and ii) all balancing energy bids are not included such as the ones from pumped hydro storage units. However, it should be noted that our market simulator makes the same market simplifications, which are thus not captured in the out-of-sample profits.

Hence, the in-sample objective value, provided by the risk-aware stochastic optimization tool, channels the understanding of both the uncertainty and risk-aware optimization models at a given risk-aversion parameter and at a specific decision stage. The purpose of our proposed ML-based module is to add an additional learning stage on top of the in-sample objective value to provide an early estimate of the out-of-sample objective value. Practically, the module is designed as a supervised ML-based regression model, which predicts online approximations of the out-of-sample objective outcomes, based on which the most suited risk-aversion parameter and, thus, the optimal risk-aware decision variables are computed.

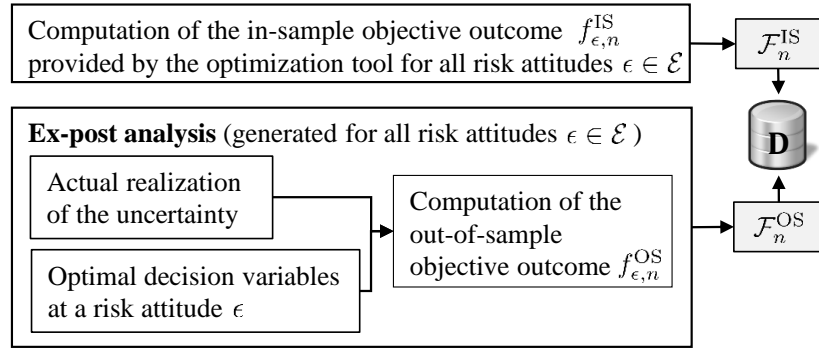
5.2.2. Training and Inference Stages of the Risk Adjustment Tool

To capture the misspecification of the uncertainty and risk-aware optimization models, the ML-based risk adjustment tool must be firstly trained on a database \mathbf{D} that maps the in-sample objective outcomes with the actual ones. As shown in Fig. 5.2, this necessary learning stage is represented by Step (A) which generates a database \mathbf{D} , i.e., the inputs $\mathcal{F}_n^{\text{IS}} = \{f_{\epsilon,n}^{\text{IS}}, \forall \epsilon \in \mathcal{E}\}$ and outputs $\mathcal{F}_n^{\text{OS}} = \{f_{\epsilon,n}^{\text{OS}}, \forall \epsilon \in \mathcal{E}\}$, on the $n = 1, \dots, N^{\mathbf{D}}$ anterior time steps, whose relationship must be learnt for different risk-aversion parameters $\epsilon \in \mathcal{E}$. The out-of-sample objective outcomes $\mathcal{F}_n^{\text{OS}}$ can be generated in an ex-post analysis for each risk-aversion level ϵ . Then, in Step (B), the objective is to optimize the parameters θ of the ML model g_θ such that we accurately map the outputs $\mathcal{F}_n^{\text{OS}}$ to the given inputs $\mathcal{F}_n^{\text{IS}}$:

$$\theta^* = \arg \min_{\theta} \sum_{n=1}^{N^{\mathbf{D}}} L(g_\theta(\mathcal{F}_n^{\text{IS}}), \mathcal{F}_n^{\text{OS}}), \quad (5.1)$$

where $L(\cdot, \cdot)$ is a user-defined loss function that quantifies how well the model

A. Construction of database \mathbf{D} ($\forall n = \{1, \dots, N^{\mathbf{D}}\}$ previous time steps)



B. Training of the ML module (ML parameters θ updated on a daily basis)




Figure 5.2.: The training procedure of the ML-based risk adjustment tool within a supervised framework.

fits the data.

In the inference stage (after the model has been trained), the ML-based module is used to predict approximations of the out-of-sample objective outcomes $\hat{\mathcal{F}}_{\text{new}}^{\text{OS}} = \{\hat{f}_{\epsilon,\text{new}}^{\text{OS}}, \forall \epsilon \in \mathcal{E}\}$ for different risk-aversion parameters $\epsilon \in \mathcal{E}$ on a new instance $\mathcal{F}_{\text{new}}^{\text{IS}} = \{f_{\epsilon,\text{new}}^{\text{IS}}, \forall \epsilon \in \mathcal{E}\}$. Then, the most suited risk-aversion parameter, and, consequently, the most optimal risk-aware decision variables can be selected based on the maximum value of the estimated out-of-sample objective outcomes. This process is showcased in Fig. 5.3.

5.2.3. Machine-Learning Models

Five Machine-Learning (ML) models are assessed and compared, including a linear model (LR), a shallow feed-forward neural network (FFNNs), random forest (RF), gradient boosted decision trees (GBDT) and the k -nearest neighbours (k -NN), for the $E = |\mathcal{E}|$ risk attitudes, where $|\mathcal{E}|$ stands for the cardinality of the set \mathcal{E} . The first four models, i.e., LR, FFNN, RF and GBDT, represent a snapshot of well-established ML techniques, as exemplified by their common use in contests such as the Global Energy Forecasting Competition [173]. The k -NN method provides another simple

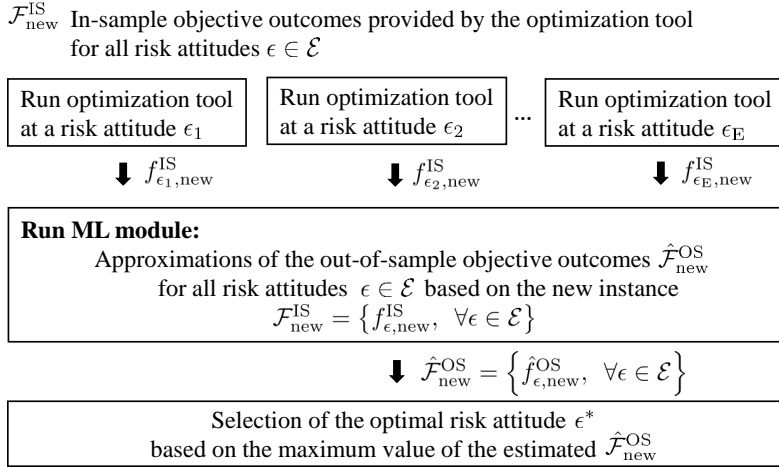


Figure 5.3.: Illustration of the automatic risk adjustment approach. In the inference stage, when a new decision stage is performed, the trained ML-based module estimates the out-of-sample objective outcomes $\hat{\mathcal{F}}_{\text{new}}^{\text{OS}}$ based on the new in-sample objective values $\mathcal{F}_{\text{new}}^{\text{IS}}$, provided by the risk-aware optimization tool for different risk attitudes $\epsilon \in \mathcal{E}$. Then, the optimal risk attitude ϵ^* is selected based on the maximum value of $\hat{\mathcal{F}}_{\text{new}}^{\text{OS}}$.

and yet competitive ML alternative, which differs from the four other methods by the fact that no model parameters θ need to be trained. In complement, the five ML models have already shown a high accuracy for approximating the objective outcomes of real-time operation processes [174], [175].

The LR model g_θ is simply expressed as $\mathbf{A}\mathcal{F}_{\text{new}}^{\text{IS}} + \mathbf{b}$, where the model parameters θ (to be optimized) are the slopes $\mathbf{A}_{(E \times E)}$ and intercepts $\mathbf{b}_{(E \times 1)}$. The loss function $L(.,.)$ is based on the least-squares criterion.

The FFNN model is the traditional architecture of neural networks, where the input information $\mathcal{F}_{\text{new}}^{\text{IS}}$ is propagated through a hidden layer containing H processing units (neurons). Each neuron consists of the application of a non-linear activation function $k_h(.,.)$, e.g., the rectified linear unit (ReLU) function, on the weighted sum of the inputs:

$$y_h = k_h \left(\sum_{\epsilon \in \mathcal{E}} w_{h\epsilon}^{\text{in}} f_{\epsilon, \text{new}}^{\text{IS}} \right) \quad \forall h \in H. \quad (5.2)$$

Then, the output vector $\hat{\mathcal{F}}_{\text{new}}^{\text{OS}}$ is given by the application of the linear function $k_o(.,.)$ on the hidden units y :

$$\hat{f}_{\epsilon, \text{new}}^{\text{OS}} = k_o \left(\sum_{h=1}^H w_{\epsilon h}^{\text{out}} y_h \right) \quad \forall \epsilon \in \mathcal{E}. \quad (5.3)$$

Using the mean square error as loss function $L(.,.)$, the backpropagation algorithm can be used to optimize the network weights $w_{\cdot}^{\{\text{in}, \text{out}\}}$.

RF and GBDT are ensemble methods based on decision trees. A decision tree \mathcal{T} is a hierarchical model combining a sequence of simple logical tests, e.g., the comparison between a numeric input and a threshold value, that recursively splits the paired dataset \mathbf{D} into distinct (smaller) subsets. Hence, the parameter θ_m of a binary decision tree at a node m is the split s_m that yields locally the *sharpest* partition of the subset \mathbf{D}_m at the node m into a left and right nodes, respectively noted m_l and m_r . The ‘sharpness’ of a split s_m at a node m can be measured by the decrease of an impurity function $i(.,.)$:

$$\Delta i(s_m) = i(m) - \frac{N^{\mathbf{D}_{m_l}}}{N^{\mathbf{D}_m}} i(m_l) - \frac{N^{\mathbf{D}_{m_r}}}{N^{\mathbf{D}_m}} i(m_r), \quad (5.4)$$

where $N^{\mathbf{D}_{\{m, m_l, m_r\}}}$ are the numbers of instances contained in $\mathbf{D}_{\{m, m_l, m_r\}}$, respectively. In a regression-based framework, $i(.,.)$ is usually derived from the

variance:

$$i(m) = \frac{1}{N^{\mathbf{D}_m}} \sum_{f_{\epsilon,n}^{\text{OS}} \in \mathbf{D}_m} \left(f_{\epsilon,n}^{\text{OS}} - \overline{f_{\epsilon}^{\text{OS}, \mathbf{D}_m}} \right)^2, \quad (5.5)$$

where $\overline{f_{\epsilon}^{\text{OS}, \mathbf{D}_m}}$ is the mean value of the outputs at the node m for a risk-aversion parameter ϵ . Such a $i(\cdot)$ results, thus, in assigning, at each terminal node z , the average value of the corresponding subset of outputs $f_{\epsilon,n}^{\text{OS}} \in \mathbf{D}_z$. Hence, the prediction of a new instance consists in (i) identifying the terminal node to which it belongs and (ii) retrieving the forecast value assigned to the corresponding node:

$$\hat{f}_{\epsilon, \text{new}}^{\text{OS}} = \mathcal{T}_{\epsilon}(\mathbf{S}_{\epsilon}, \mathcal{F}_{\text{new}}^{\text{IS}}) \quad \forall \epsilon \in \mathcal{E}. \quad (5.6)$$

where \mathbf{S}_{ϵ} define the splits of the decision tree \mathcal{T}_{ϵ} at a risk-aversion parameter ϵ .

In RF, an ensemble of N^{RF} decision trees are independently grown using the aforementioned approach, with the particularity that each split of each tree is constructed based on a random subsample of the data set and a random subset of features, allowing the reduction of the variance of the entire model. A new prediction is obtained by averaging the outcomes of each decision tree:

$$\hat{f}_{\epsilon, \text{new}}^{\text{OS}} = \frac{1}{N^{\text{RF}}} \sum_{i=1}^{N^{\text{RF}}} \mathcal{T}_{\epsilon,i}(\mathbf{S}_{\epsilon,i}, \mathcal{F}_{\text{new}}^{\text{IS}}) \quad \forall \epsilon \in \mathcal{E}. \quad (5.7)$$

As for the GBDT method, the trees are not built independently but rather in an additive fashion. For the p -th iteration, the GBDT model g_{θ} is written as:

$$g_{\theta}^p(\mathcal{F}_{\text{new}}^{\text{IS}}) = g_{\theta}^{p-1}(\mathcal{F}_{\text{new}}^{\text{IS}}) + \alpha \mathcal{T}^p(\mathcal{F}_{\text{new}}^{\text{IS}}), \quad (5.8)$$

where α is the learning rate. At each stage, the additional tree \mathcal{T}^p is updated to minimize the residuals of the $p - 1$ previously generated trees:

$$\mathcal{T}^p = \arg \min_{\mathcal{T}} \sum_{n=1}^{N^{\mathbf{D}}} L\left(f_{\epsilon,n}^{\text{OS}}, g_{\theta}^{p-1}(\mathcal{F}_n^{\text{IS}}) + \mathcal{T}(\mathcal{F}_n^{\text{IS}})\right), \quad (5.9)$$

where $L(\cdot, \cdot)$ is the mean square error function, and $N^{\mathbf{D}}$ is the number of instances contained in the database \mathbf{D} . The overall approach is similar to the gradient descent algorithm in which the added tree \mathcal{T}^p is optimized to leverage the prediction errors of its predecessors.

The k -NN algorithm relies on the concept of learning by analogy: based on the creation of a paired dataset, the prediction of a new instance is carried out on the average value of the k nearest neighbours in the dataset where the closeness condition is determined based on a distance metric (e.g., the euclidean distance). For a new instance, the k -NN estimate is given by:

$$\hat{\mathcal{F}}_{\text{new}}^{\text{OS}} = g(\mathcal{F}_{\text{new}}^{\text{IS}}) = \frac{1}{k} \sum_{n=1}^{N^{\text{D}}} x_n \mathcal{F}_n^{\text{OS}}, \quad (5.10)$$

where $x_n \in \{0, 1\}$ depending on whether or not $\mathcal{F}_n^{\text{IS}}$ is among the k -nearest neighbours of $\mathcal{F}_{\text{new}}^{\text{IS}}$ and N^{D} is the number of instances within the database \mathbf{D} .

In practice, new information is continuously revealed and the training procedure must be updated over time. In this context, an appealing feature of k -NN is that it seamlessly supports online updates [176]. Its prediction performance is simply improved by adding new instances to the dataset, while the other ML models need the additional Step **(B)** to calculate their optimal ML parameters θ^* . In our case study focusing on single price imbalance settlement markets, the parameters θ^* are updated on a daily basis resulting in a balance between computational burden and model precision.

In general, each of the ML models are characterized by hyper-parameters, which are (task-dependent) parameters reflecting the complexity of the model, e.g., the number of hidden units of neural networks, the number of basic trees in RF or the number of K neighbours. These values are estimated using the hyperparametrization approach shown in Chapter 3.

5.3. Case Study

Similarly to Section 4.9 in Chapter 4, the case study leverages the probabilistic predictions of the system imbalance obtained via the Sequence-to-sequence recurrent neural model. The prediction model is trained using historical data from 2014 until end of December 2016, and is stabilized with early stopping using the year 2017 as a validation set to avoid overfitting. Then, the probabilistic predictions are inferred at the start of each quarter hour of the test set, i.e., the year 2018. At each time step, the feasible region of the market actor portfolio Π_{UL} is constrained by (upward and downward) 120 MW power limits. The cost parameters C^+ and C^- are respectively set to 50 and 30 €/MWh ensuring a consistent imbalance position regarding the imbalance price. We consider 11 risk attitudes for the actor corresponding to

the following $\epsilon \in \{0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 0.98\}$ risk-aversion parameters. It should be recalled that the risk-aversion parameters are i) the prediction interval $\epsilon = 1 - 2q$ characterizing the size of the uncertainty set in robust optimization (4.9), and ii) the confidence level ϵ defining the $(1 - \epsilon) \times 100\%$ worst-case scenarios in the CVaR-based stochastic program (4.6). The case study investigates whether an automatic risk adjustment of the risk-aware decision-support tools can provide an ex-post economic gain.

To emphasize the learning ability of the ML tools supporting the selection of a risk policy, the models are trained only on the data available during the year 2018. The first half of January is used to tune the hyperparameters of the models using a random search embedded within a cross validation scheme, and the latter are then recalibrated on a daily basis, using the new information revealed over time. The final five ML regression models (along with their search spaces) are:

- An LR model.
- An 1-MLP model with $H = 20$ neurons using early stopping, and rectified linear units (ReLU) as activation functions. The search range of H was $\{10, 20, 50, 100\}$.
- An RF model with $N^{RF} = 100$ and a maximum depth of 5. We have varied the maximum depth between $\{3, 5, 8\}$.
- An GBDT model, in which $\alpha = 0.05$, the maximum depth is 3, and the number of iterations is determined by using early stopping. The search range of α was $\{0.3, 0.2, 0.1, 0.05\}$, and we also varied the maximum depth between $\{3, 5, 8\}$.
- An k -NN model with $k = 1000$, where the search range of k was $\{1, 100, 500, 1000, 1250\}$.

Interestingly, it should be noted that the hyperparameter k of the k -NN model can be naturally interpreted. If $k = 1$, the selection of the risk policy is only based on the ex-post economic performance of the closest instance within the database (where the closeness condition is computed based on the in-sample objective values using the euclidean distance). On the other hand, if k is set to the number of instances within the database, the risk policy's selection of the k -NN method would coincide with the one obtained using the traditional offline out-of-sample analysis (which is performed at this specific time step).

The results are computed over the period spanning from 15 January until the end of December 2018. The notation RO-Q stands for the robust optimization, while SP-CVaR denotes the CVaR-based stochastic optimization. The time step Δt is equal to an imbalance settlement period, i.e., 15 minutes. The performance of the proposed methodology is principally evaluated through one indicator, which is averaged over all the quarter hours of the test set: the actual profit $f^{OS}(\cdot) = (\lambda^{SI} - C^+)e^{imb,+} - (\lambda^{SI} - C^-)e^{imb,-}$ obtained during the ex-post analysis. All the experiments have been conducted on an Intel Core i7-8850H CPU running at 2.60 GHz and with 16.0 GB of RAM, and coded in Python 3.6. The probabilistic forecasting model and the 5 ML models are implemented using the Scikit-learn and Keras packages. The MILP formulations for the risk-aware optimization models are written using the Pulp package and solved using the Gurobi 8.1.1 solver.

Table 5.1 shows the evolution of the ex-post profits achieved by the different ML techniques over several time periods of the test set for RO-Q and SP-CVaR. These results are put into perspective with i) Online: the ideal choice of the risk attitude at each quarter hour which is determined ex-post when the actual conditions are revealed and ii) Offline: the single choice of a risk-aversion parameter ϵ based on the ex-post economic performance of the different risk-aversion parameters \mathcal{E} (from risk-seeking $\epsilon = 0$ towards risk-averse $\epsilon = 0.98$) during the first half of January. Over the 1-year period, the (omniscient) risk policies Online for each optimisation methodologies allow to improve the real revenues by over 121% and 57 % compared to their counterparts Offline, stressing the relevance of adopting a dynamic risk attitude. As expected, the optimal risk policy Online for robust approach leads to greater ex-post profits than the SP-CVaR. This gap can be explained by the difference in the risk attitude at $\epsilon = 0$. At this risk-aversion parameter, the robust approach is purely deterministic and is extremely risk-seeking which leads to huge rewards in case of perfect information, while SP-CVaR is only risk-neutral and still takes into account extreme scenarios in its decision-making, preventing to fully leverage the added value of a perfect forecast.

Overall, the proposed online selection of the risk attitudes (guided by the best ML models) respectively improves the ex-post economic performance by 12.3% and 4.8% for RO-Q and SP-CVaR in comparison with their counterparts Offline over the 1-year period, which highlights the added value of better informing stochastic optimization tools with tailored risk-aversion parameters. In addition, the ability of ML to dynamically select the optimal risk parameter is analyzed in Fig. 5.4, by showing the frequency at which each risk-aversion

Table 5.1.: Rolling average of the out-of-sample profits $f^{OS}(\cdot)$ over different time periods. The percentage values indicate the variation of the out-of-sample profit with respect to the Offline strategy.

[€/Δt]	Offline	LR	FFNN	RF	GBDT	k-NN	Online
RO-Q	$\epsilon = 0.7$						
1 day	127.6	107.8 (▼15.5%)	51.1 (▼60%)	99 (▼22.4%)	102.4 (▼19.7%)	124.6 (▼2.3%)	317.8 (▲149%)
1 week	255	247.9 (▼2.8%)	271.6 (▲6.5%)	282.4 (▲10.7%)	255.9 (▲0.3%)	234.6 (▼8%)	633.6 (▲148.5%)
1 month	198.2	205.5 (▲3.7%)	202.1 (▲2%)	223.1 (▲12.6%)	201.8 (▲5%)	208 (▲5.1%)	528.4 (▲166.6%)
3 months	340.5	362.4 (▲6.4%)	349.2 (▲2.6%)	377 (▲10.7%)	367.5 (▲7.9%)	360.9 (▲6%)	696.4 (▲104.5%)
6 months	343.1	343.1 (▲0%)	341.5 (▼0.5%)	366.5 (▲6.8%)	360.4 (▲5%)	361.6 (▲5.4%)	693.3 (▲102.1%)
1 year	281.2	288.3 (▲2.5%)	294.1 (▲4.6%)	312.9 (▲11.3%)	312 (▲11%)	315.7 (▲12.3%)	622.4 (▲121.3%)
SP-CVaR	$\epsilon = 0.3$						
1 day	135.2	112.1 (▼17%)	131 (▼3.1%)	105.3 (▼22.1%)	127.6 (▼5.6%)	123.12 (▼8.9%)	289 (▲113.8%)
1 week	323.5	319.2 (▼1.3%)	312.4 (▼3.4%)	297.1 (▼8.1%)	286.6 (▼11.4%)	311.7 (▼3.7%)	594.1 (▲83.6%)
1 month	238.8	242 (▲1.3%)	240 (▲0.5%)	234.2 (▼1.9%)	229 (▼4.1%)	245.7 (▲2.9%)	496.9 (▲108.1%)
3 months	396.5	410.9 (▲3.6%)	398.5 (▲0.5%)	408.3 (▲3%)	396.9 (▲0.1%)	416.9 (▲5.1%)	672.3 (▲69.6%)
6 months	375.1	393.2 (▲4.8%)	380 (▲1.3%)	392.6 (▲4.7%)	382.7 (▲2%)	398.8 (▲6.3%)	662.1 (▲76.5%)
1 year	327	338.1 (▲3.4%)	327.9 (▲0.3%)	337.3 (▲3.1%)	333.3 (▲1.9%)	342.7 (▲4.8%)	515.7 (▲57.7%)

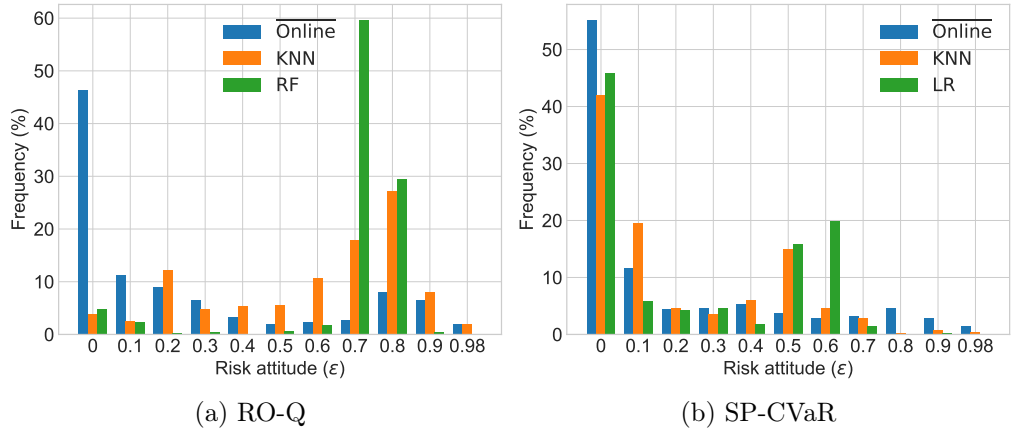


Figure 5.4.: Histograms of the risk-aversion parameters selected over the 1-year period.

parameter was chosen over the test set. For clarity, only the two best ML models are represented, i.e., $\{k\text{-NN}, \text{RF}\}$ for RO-Q, and $\{k\text{-NN}, \text{LR}\}$ for SP-CVaR.


Regarding RO approaches, the 1 year-based results in Table 5.1 show that an online support in the construction of the uncertainty set is a key element to fully leverage their potential. Indeed, we see that each ML model allows outperforming the economic gains of the strategy Offline, emphasizing the importance of the size of the uncertainty set and its effect on the economic performance for robust-based optimization formulations. Logically, the simplest model LR gives less insights about the selection of the risk-aversion parameter than more advanced ML techniques. In particular, the models LR and FFNN show lower economic performances than the ensemble methods (RF and GBDT). Interestingly, the $k\text{-NN}$ technique, which is simple and intuitive, shows a high suitability for our application as it outperforms the other ML techniques over the 1-year period. Additionally, Fig. 5.4a demonstrates that both ML techniques i) select predominantly (45% and 89% of the time for respectively $k\text{-NN}$ and RF) the optimal risk policies Offline at $\epsilon = \{0.7, 0.8\}$ and ii) timely deviate 16% ($k\text{-NN}$) and 7.5% (RF) of the time towards riskier strategies at $\epsilon = \{0, 0.1, 0.2, 0.3\}$. These distributions explain the gap in ex-post profit with respect to the optimal (omniscient) strategy Online, which requires a greedy approach that adopts more than 45% of the time the most risky strategy.

Concerning SP-CVaR, none of the ML-based strategies provide worse ex-post profit than the offline one over the period of one year, but their


added value in terms of ex-post profits is lower in comparison with the robust case. Surprisingly, the simple LR model provides a better ex-post performance than the other (more complex) parametric ML models (1-MLP, RF and GBDT), while k -NN devises the most optimal online risk strategy. Fig. 5.4b) gives a first rationale behind such results by showing the vision of the risk strategies supported by ML: they adopt regularly the risk-neutral strategy (more than 40% of the time), while sparsely selecting risk-averse strategies at $\epsilon = \{0.5, 0.6\}$. This kind of risk management, recommended by the optimal one Online, appears to be more challenging to implement for the more advanced ML models. Indeed, the latter are inclined to adopt a more conservative behaviour, which penalize them more severely in the SP-CVaR case.

Regarding the evolution over time of the RO-Q and SP-CVaR ex-post profits, it can be seen that the ML-based risk strategies have the ability to learn and capture the adequate risk attitude rapidly, which is reflected by the positive gains after only a day on the field. In the same vein, for both optimization tools, the machine learning approaches exhibit better economic performance than the offline one after only one month. However, the gap between ML-based strategies and the approach Offline does not widen after the first month of use. It tends to show that specific calibration methods for ML models have to be developed to take full advantage of the new data that are constantly revealed over time [177]. More particularly, the k -NN method is less affected by the calibration method. The results of Table 1 indicate that the k -NN method is emerging over time as the most suited method for supporting the automatic risk adjustment strategy of a market player.

In this line, Fig. 5.5 compares the histograms of the out-of-sample profits between the Offline, k -NN and Online strategies over the 1-year period. It should be noted that the null out-of-sample profits are not represented in both histograms. Concerning the RO-Q approach, the out-of-sample profits are null in 43%, 31% and 30% of the time for Offline, k -NN and Online, respectively. For the SP-CVaR approach, the out-of-sample profits are null in 39%, 21% and 29% of the time for Offline, k -NN and Online, respectively. In both histograms, the out-of-sample profits mostly lie between -500€ and 500€. We can observe that the k -NN approach generates more negative revenues than the Offline one between [-500,0], while it yields higher gains compared to the Offline one in the range [0,500]. Hence, the k -NN approach improves its average ex-post profit in comparison with the Offline one by providing more regularly real-time balancing services, which allows capturing higher opportunity profits but with an increased risk exposure. As expected, the Online approach systematically results in null or positive revenues.



(a) RO-Q



(b) SP-CVaR

Figure 5.5.: Histograms of the out-of-sample profits performed at each quarter hour over the 1-year period.

Table 5.2.: TSO Balancing actions and cost for different risk strategies (the percentage values indicate the variation of the balancing actions and cost with respect to the non-participation of a market player).

	$\mathbb{E}(V^{\text{act}})$ [MW/ Δt]	$\mathbb{E}(C^{\text{os}})$ [€/Δt]
No Participation	24.36	1350.11
RO-Q with k -NN	19.22 (▼21.1%)	782.34 (▼42%)
SP-CVaR with k -NN	19.39 (▼20.4%)	717.83 (▼46.8%)

The ML-based approach is more computational intensive than Offline strategies as it requires the prior computations of the in-sample profits $\mathcal{F}^{\text{IS}} = \{f_\epsilon^{\text{IS}}, \forall \epsilon\}$ as inputs. In our case study, the averaged computation time of an in-sample profit f_ϵ^{IS} is 0.02s for the robust optimization framework and 1.1s for the CVaR-based stochastic optimization framework with a duality gap of 1% imposed. Concerning the inference time of the LR, 1-MLP, RF, GBDT and k -NN methods, their averaged time for outputting the most optimal risk attitude are around 0.05ms, 27ms, 112ms, 5ms and 1ms, respectively. Overall, the prior computations of the in-sample profits \mathcal{F}^{IS} are the most time consuming, but this issue can be alleviated through parallel computing. However, there still exists a gap in computation times between RO-Q and SP-CVaR, e.g., RO-Q is 50 times faster than SP-CVaR in our application. If the computation time is a hard constraint, the RO-Q approach guided by the k -NN model provides a viable alternative to the Offline SP-CVaR approach. This allows reducing the ex-post profits differential between RO-Q and Offline SP-CVaR from 16.3% (Offline RO-Q) to 3.6% (RO-Q with k -NN model) over the entire year.

Finally, the impact of the imbalances of the Balance Responsible Party (BRP) on the performance of the TSO balancing dispatch procedure is studied in Table 5.2. To that end, the cost $C^{\text{os}} = \sum_{r^+ \in R^+} \Lambda_{r^+} s_{r^+} + \sum_{r^- \in R^-} \Lambda_{r^-} s_{r^-}$ and the activated energy $V^{\text{act}} = \sum_{r^+ \in R^+} s_{r^+} + \sum_{r^- \in R^-} s_{r^-}$ of each quarter hour are retrieved from the ex-post analysis and are averaged over the test set. Results show an average drop of 20% and 40% for, respectively, the averaged activation of balancing energy and the costs, thereby reducing the TSO's corrective actions at the real-time balancing stage.

5.4. Conclusion

This Chapter leverages the self-learning abilities of ML techniques to dynamically and preemptively adjust the risk policy of a Balance Responsible Party (BRP) based on the current state of its expected market outcomes. Its risk policy is progressively updated and improved based on past trading sessions. The effectiveness of the proposed automatic risk-adjustment tool is illustrated in a detailed case study using data from the Belgian power system with a wide range of competitive ML-based techniques (i.e., linear regression, neural networks, tree-based ensemble methods and k -nearest neighbours). The automatic risk-adjusted policies, considering both robust optimization and CVaR-based stochastic programs, are then compared with traditional time-invariant risk policies, highlighting the economic potential of adopting a dynamic risk policy.

Indeed, extensive numerical analyses using real-world market data from the Belgian power system over one year demonstrate the ability of the proposed approach to achieve efficient online risk-adjusted strategies for robust-based and CVaR-based stochastic optimizations. More particularly, the k -NN technique has been identified as a suited ML candidate to support these risk-aware optimization methods for preemptively devising the risk attitude. In this line, both RO-Q and SP-CVaR approaches guided by the k -NN model have presented promising results as this had led to a respective increase of 12.3% and 4.8% in the ex-post profits compared with their offline risk policy-based counterparts. Besides, the implementation of our theoretical models on actual electricity market settings corroborate the key goal of the single price imbalance settlement mechanism, by reducing the system imbalance, and consequently limiting corrective actions at the real-time balancing stage. Indeed, the obtained results show that the BRP increases its operating profit, while the imbalance of the power system is reduced.

In that research direction of improving the interaction of the user with automated data-driven decision-support tool, the following Chapter is devoted on combining the predictive power of neural models with interpretability. The analysis of interpretable forecasting outcomes may allow to deliver sanity checks for the designer and/or user for increasing its confidence on the associated forecaster.

Chapter Publication

- **J. Bottieau**, K. Bruninx, A. Sanjab, Z. De Grève, F. Vallée and J-F. Toubéau, "Automatic Risk Adjustment for Profit Maximization in

*Chapter 5. Automatic Risk-Adjusted Provision of Real-Time Balancing
Services*

Renewable Dominated Short-Term Electricity Markets," in ITEES, vol.
3, issue 12, 2021

CHAPTER 6.

Towards Interpretable Probabilistic Forecasting Using Neural Networks

Chapter 3 is essentially focused on i) better capturing the non-linear and uncertain behavior of the system imbalance, and ii) fully exploiting the temporal information contained in both past observed and future known input data. This has led to the attention-based sequence-to-sequence recurrent neural model, which has demonstrated high-quality performance compared to other competitive time series forecasting models. However, even an highly accurate predictive model may face barriers in terms of acceptability among the users' community if it behaves as a black box, where its inner prediction process is hardly understandable. The neural models are particularly prone to that phenomenon, as their underlying reasoning is more complex to extract than simpler, readily interpretable models (such as the AutoRegressive moving-average (ARMA) model). In this line, combining the predictive power of deep neural models with interpretable features has attracted a high interest within the machine learning community, especially in the computer vision and natural language processing fields [178], [179]. However, there is still a lack of dedicated research for time series applications [180].

This Chapter falls within this research line for deep learning-based time series forecasting methods, aiming at outputting accurate predictions, while identifying the most important input features of the model and their interaction. In this Chapter, we focus on directly forecasting the imbalance prices, without passing through the system imbalance prediction and the balancing energy market-clearing stages. Indeed, being a signal characterized by sharp regime switching, predicting directly the imbalance prices would allow to bring to light distinct temporal patterns of the input signals depending on the predicted price regime. The structure of this Chapter is organized as follows. A brief status of interpretability in neural models is presented in Section 6.1, while Section 6.2

presents a review on imbalance price forecasting. In Section 6.4, the benchmark methods are presented. Then, the case study and the evaluation of the proposed forecasting strategy, both in terms of performance and interpretability perspectives, are discussed in Section 6.5. Finally, Section 6.6 concludes the Chapter.

6.1. Interpretability in Neural Models

The notion of interpretability have received a great resurgence of interest lately for opening black-box deep neural models. Although this notion has been intensively explored in the literature, no clear consensus on the definition of interpretability has been reached [181], [182]. Following [183], interpretability is the ability to provide explanations in understandable terms to a human, where: i) explanations, ideally, should be logical decision rules (if-then rules). However, in practice, the explanations are usually reduced to only highlight some key input drivers of the model [179]. ii) Understandable terms, which indicates that the explanations should rely on accepted knowledge related to the task. Based on this definition, this Chapter aims at developing an interpretable neural model able to provide insights of its inner input-output mappings to the designer and/or end-user.

In the same vein, measuring interpretability for time series models is also not straightforward, and is still under research since no ground truth can be used for benchmarking purposes. One way for bypassing that is to generate synthetic time series data, where multiple feature-time interactions of gradual complexity can be artificially designed – see, e.g., [180]. On the other hand, for measuring interpretability on real-life time series, two different aspects are commonly benchmarked: i) precision (whether all the identified features are relevant), and ii) recall (whether all the relevant features are correctly identified). In this Chapter, interpretability is rather used as a tool for checking whether the model has captured a meaningful causality between the labeled important input features and the predicted outputs based on human knowledge.

Interpretability is challenging for energy forecasting as it is characterized by a high space-time input domain (e.g., the weather conditions at different locations of the grid). Besides, the visualization of raw time series contains less expressive power than images or texts for the non-expert user. Both aspects have rendered the development of interpretable energy forecasting models to be lagging behind compared to other fields such as computer vision or natural language processing. Indeed, energy forecasting models are essentially benchmarked based on accuracy metrics, which favors the trend towards overly complex

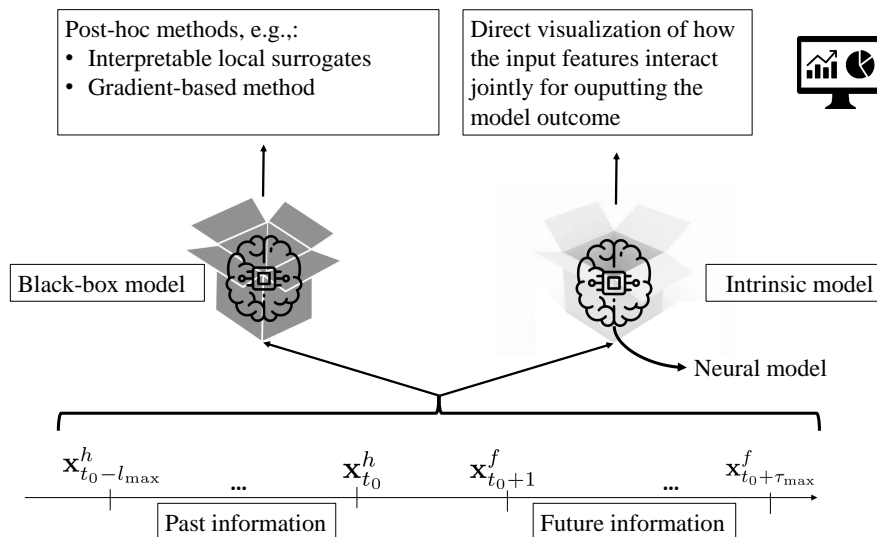


Figure 6.1.: Interpretability methods for neural models

black-box models. Yet, developing an interpretable model can be beneficial for i) the designer, who can check that the model does not exploit artifacts in the data, and ii) the decision maker, who will be better equipped for making its decision. Indeed, interpretability may mitigate the risk of adoption of machine learning models for a reluctant decision maker. Besides, the decision maker can confront its own business logic and intuitions with the inner prediction process of the model, which may provide novel insights about the forecasting application. Finally, an increasing number of legal requirements on machine learning models (e.g., the General Data Protection Regulation 2016/679, GDPR) push the business entities towards the adoption of interpretable models.

Two distinct methods can be applied for interpretability in neural models (see Fig. 6.1): i) post-hoc methods, which consist in analyzing an *already trained* (black-box) model (e.g., interpretable local surrogates or gradient-based methods) [184]. However, such methods are limited when applied to multivariate time series as they typically do not consider the temporal dependencies between features [180]. ii) Intrinsic methods, in which the architecture of the neural model is directly designed with interpretable components [185], [186].

In post-hoc methods, the surrogate-based approach provides local interpretable insights of an outcome (of a black-box model) by learning a simpler model based on crafted instances close to the associated input (i.e., small perturbations are added to the input, which are then labeled by the original

model). The most well-known method is Local Interpretable Model-Agnostic Explanations (LIME) developed in [187], which learns a linear model on top of these perturbed samples for locally attributing an importance value at each input feature. In the same vein, the Shapley additive explanation method is developed in [188], which calculates the average contribution of each feature by comparing the model outcome for different permutations of input feature. On the other hand, gradient-based methods can also be used to provide insights of the inner working of neural models. Hence, by computing the gradients of the network outcome with respect to the input, the obtained values can represent the sensitivity of the model output with respect to each specific feature [189]. For instance, in computer vision, the so-called ‘saliency’ maps (i.e., a 2-dimensional map attributing an importance value for each pixel of an input picture) rely on gradient-based methods, and allow checking if the model vision agree with human intuition [183].

Besides post-hoc methods, research efforts have also covered intrinsic methods, which attempts at imposing interpretability pathways during the network training process. First, regularization terms can be directly added to the loss function, for imposing, e.g., sparsity in the neural model (where only a few input features can interact jointly) [190]. Beyond tweaking the loss function, the attention mechanism can be viewed as an architectural pathway, guiding the neural model on how it should process temporal information. Hence, the attention weights that are computed internally by the model can thus present some inherent interpretable characteristics. In addition to the attention mechanism, this Chapter augments the neural model by adding feed-forward neural layers that are designed to endogenously quantify the relative importance between input features at each time step before being processed by the model.

6.2. Probabilistic Forecasting of Imbalance Prices

So far in the report, the (real-time) imbalance prices were obtained by i) forecasting the system imbalance (Chapter 3), followed by ii) the computation of the market-clearing proxy of the balancing energy market (Chapter 4). One drawback of this two-step approach is the implicit adoption of the simplifying market hypotheses made by the Transmission System Operator (TSO) when computing the merit order proxies of the balancing energy bids. This inevitably limits the accuracy of the predicted imbalance prices. In this context, this Chapter proposes a more straightforward approach, where the (real-time) imbalance prices are directly predicted. Yet, the prediction of

such real-time prices is challenging mainly due to two fundamental causes: i) the signal exhibits a regime-switching behavior, where it flips from low- and high-price regimes depending on whether the power system is in surplus or shortage of generation [191], and ii) price spikes occur more frequently due to the market's small size and vulnerability to unexpected changes in operating conditions, e.g., outages or congestion of transmission lines [40]. It should be recalled that, in European markets, the term 'real-time electricity prices' may refer to either imbalance or balancing prices, which arises from the intrinsic segmentation between the energy and balancing markets (see Chapter 2). Hence, balancing prices remunerate the Balancing Service Providers (BSPs) for the actual activation of balancing reserves (e.g., an automatic frequency restoration product), whereas imbalance prices monetize any real-time energy deviations of Balance Responsible Parties (BRPs) from their position in energy markets. Both prices are connected since the imbalance price is based on the marginal price of the activated balancing energy. In contrast, the real-time electricity prices in US-styled pools are defined using a locational marginal pricing market, wherein energy deviations and operating reserves are settled at a unique price for each electrical node. In this Chapter, without loss of generality and for consistency reason with the rest of the report, we focus on the probabilistic prediction of (Belgian) imbalance prices.

Despite the inherent difficulty in predicting real-time electricity prices, the literature is still scarce compared to, e.g., the day-ahead electricity prices – see, e.g., [192], [193] and references therein. Markov regime-switching models have been proposed as natural candidates for capturing the real-time prices [191], [194], [195]. More specifically, Olsson and Söder present a Markov-switching seasonal auto-regressive moving average model (M-SARIMA) in [194], while they investigate the introduction of exogenous variables using non-linear time series models in [196]. In the same vein, Dimoukias et al. apply a hidden Markov model for modeling Nordic balancing prices [195], while Bunn et al. analyze the predictability of British balancing prices using Markov switching dynamic regression models [191]. Although these models have proved good properties, one issue arises with the computation of the transition probabilities, which relies on the well-known Markov property stating that the expected future state of the process depends only on its present state. To avoid this issue, Klæboe et al. perform a benchmark analysis of time-series based forecasting models for Nordic balancing prices, which tends to show that embedding the balancing state information (i.e., actual imbalance volumes in the system) in the forecasting models provide sharper interval forecasts [197]. In that way, a SARIMA model based on the activated balancing volume is proposed in [198], and a Holt-Winters model conditioned by the sign of the net imbalance

volume is developed in [199]. The importance of the net system imbalance volume is further highlighted in [200], which shows that this covariate has the highest explanatory power when used with tree-based ensemble methods. Overall, the literature shows that capturing both the price-regime switching behavior and spikes of real-time electricity prices is a non-trivial task [201], [202].

In this Chapter, we propose to use a Transformer model [203] for forecasting the real-time electricity prices, whose architecture is augmented for fostering its intrinsic interpretability. Transformers tend to become the novel state-of-the-art neural model in various tasks such as, e.g., natural language processing [204]. By relying on attention mechanisms solely computed via feed-forward neural networks, the proposed model is here designed to capture distinct temporal patterns of the input signal depending on the predicted price regime. In addition, the model is augmented with subnetworks able to provide direct insights on the relative importance of each individual input feature [186]. Finally, as deep learning models are known to be difficult to optimize and require careful tuning of hyper-parameters, new regularization techniques are leveraged to improve both the convergence and performance of the proposed model [205].

6.3. Interpretable Transformer Model

The model is designed for generating multi-horizon probabilistic forecasts of the real-time price λ^{RT} for each imbalance settlement period (ISP):

$$p\left(\lambda_{t_0+1}^{\text{RT}}, \dots, \lambda_{t_0+\tau_{\max}}^{\text{RT}} \mid \mathbf{x}_{t_0-l_{\max}}^h, \dots, \mathbf{x}_{t_0}^h, \mathbf{x}_{t_0+1}^f, \dots, \mathbf{x}_{t_0+\tau_{\max}}^f\right) \quad (6.1)$$

where t_0 is the forecast creation time, l_{\max}, τ_{\max} are respectively indices determining the number of look-back and look-ahead ISP, $\mathbf{x}^h \in \mathbb{R}^{m_h}$ are time series observed, and $\mathbf{x}^f \in \mathbb{R}^{m_f}$ are future information, e.g., the prices cleared at the day-ahead stage or calendar information, already known over the prediction horizon.

From the literature review in [206] and insights of Chapter 3, three major trends can be identified in the development of neural models: i) deeper architectures, when adequately designed, increase the ability of the network to extract meaningful representation from the raw data, ii) convolutional neural networks or recurrent neural networks – especially, the Long Short Term Memory (LSTM) – are efficient in learning local spatio-temporal relationships, and iii) attention mechanisms, which grant the model direct access to information on specific time steps, enable an improved representation of long-term dependencies.

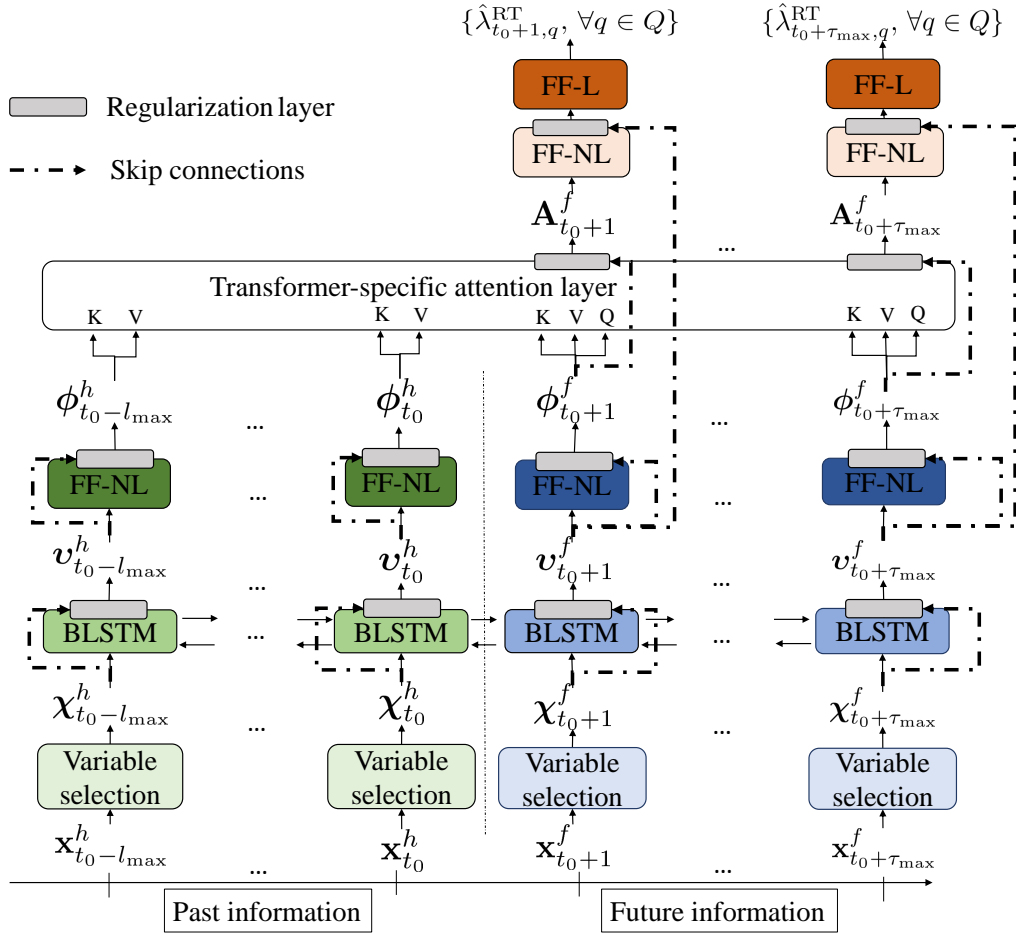


Figure 6.2.: The transformer-based model.

In light of these recent advances, we propose a transformer-based model, which pursues high-quality probabilistic predictions of real-time electricity prices, while attaining interpretable insights. The overall model is depicted in Fig. 6.2. Note that layers in the same color share the same weights. In addition, the notations FF-NL and FF-L stand for feed-forward networks using respectively non-linear and linear activation functions, while BLSTM refers to a bi-directional LSTM network. At the early stage, the m_h look-back observed inputs and the m_f look-ahead known inputs are respectively processed by two distinct variable selection subnetworks, which act as an interpretable filter that allows the model to disregard any irrelevant inputs (Subsection 6.3.2). The selected inputs are then handled by bi-directional

long short-term memory networks, where both backward and forward time correlations are locally captured (Subsection 6.3.3), followed by a feed-forward network that computes an additional non-linear mapping if required. For each time step of the prediction horizon, the Transformer-specific attention layer selectively identifies the most salient past and future contextual information over the conditioning range $[t_0 - l_{\max}, t_0 + \tau_{\max}]$ in a single vector representation (Subsection 6.3.4). Finally, based on this condensed representation, a direct multi-horizon strategy is applied, which consists in outputting in one pass the real-time price's q -quantiles $\{\hat{\lambda}_{t_0+\tau, q}^{\text{RT}}, \forall q \in Q\}$ through two successive non-linear and linear mappings. This strategy avoids error accumulation (which is common in fully recurrent models) by alleviating the need of recursively feeding the previously predicted target, while fully making use of the parallel abilities of hardware such as GPUs. In addition, throughout the model, we repeatedly used regularization layers and skip connections (Subsection 6.3.5) to control the complexity of the model and facilitating its training. All the layers are trained in an end-to-end fashion, i.e., all layers are jointly trained, which guarantees the consistency of the framework.

6.3.1. Feed-Forward Networks

Feed-forward networks are used for either transforming a n -dimensional input vector into a d -dimensional vector or applying additional linear and non-linear mappings. Let $\mathbf{x}^{\text{in}} \in \mathbb{R}^n$ be the input vector. The linear mapping of an FF-L layer is defined as:

$$\mathbf{x}^{\text{out}} = \mathbf{x}^{\text{in}}\mathbf{W}_1 + \mathbf{b}_1 \quad (6.2)$$

where $\mathbf{x}^{\text{out}} \in \mathbb{R}^d$ is the d -dimensional output vector, and $\mathbf{W}_1 \in \mathbb{R}^{n \times d}$ and $\mathbf{b}_1 \in \mathbb{R}^d$ are parameters to be trained.

An FF-NL layer consists of two linear transformations, with a non-linear activation function in between:

$$\mathbf{x}^{\text{out}} = f^{\text{ELU}}(\mathbf{x}^{\text{in}}\mathbf{W}_2 + \mathbf{b}_2)\mathbf{W}_3 + \mathbf{b}_3 \quad (6.3)$$

where $\mathbf{W}_2 \in \mathbb{R}^{n \times d}$, $\mathbf{W}_3 \in \mathbb{R}^{d \times d}$, $\mathbf{b}_2 \in \mathbb{R}^d$ and $\mathbf{b}_3 \in \mathbb{R}^d$ are parameters to be trained, and f^{elu} is the exponential linear unit (ELU) activation function, which acts as an identity function for positive values and gets saturated for negative ones [207].

6.3.2. Variable Selection Layer

Similarly to Chapter 3, besides calendar information, we have at our disposal $m_h = 14$ historical covariates $\mathbf{x}_{:t_0+1}^h$ and $m_f = 15$ known future information $\mathbf{x}_{t_0+1}^f$. These inputs are:

- the imbalance price ($\lambda^{h,\text{RT}} \in \mathbb{R}^1$).
- the net activated volume of balancing reserves ($\text{NRV}^h \in \mathbb{R}^1$).
- the upward and downward balancing prices ($\lambda^{h,\text{bal.}} \in \mathbb{R}^2$).
- the physical cross-border energy flows with France and Netherlands ($\phi^h \in \mathbb{R}^2$).
- the produced and forecasted wind and photovoltaic powers with their associated installed capacities ($P^{\{h,f\},\text{renew.}} \in \mathbb{R}^4$).
- the produced and scheduled powers of conventional generators ($P^{\{h,f\},\text{conv.}} \in \mathbb{R}^3$), composed of pump-hydro, gas and nuclear units.
- the measured and forecasted electrical load of the grid ($L^{\{h,f\}} \in \mathbb{R}^1$).
- the day-ahead electricity prices ($\lambda^{f,\text{DA}} \in \mathbb{R}^1$).
- the merit order proxies of operational balancing prices, i.e., the TSO expected prices corresponding to different volumes of activated reserves $\{-600, -300, -100, 100, 300, 600\}$ MW ($\lambda^{f,\text{bal.}} \in \mathbb{R}^6$).

It should be noted that all the continuous inputs are firstly min-max normalized between $[-1, 1]$. The calendar information ($\mathbf{x}^{\{h,f\},\text{cal.}} \in \mathbb{R}^6$) are categorical variables characterizing working days, the day of the week, the hour, the quarter hour, the month and the absolute position of the time step. Overall, the set of historical covariates $\mathbf{x}_{:t_0+1}^h$ is composed of $\{\lambda^{h,\text{RT}}, \text{NRV}^h, \lambda^{h,\text{bal.}}, \phi^h, L^h, P^{h,\text{renew.}}, P^{h,\text{conv.}}, \mathbf{x}^{h,\text{cal.}}\}$, while the set of future known information $\mathbf{x}_{t_0+1}^f$ contains $\{L^f, P^{f,\text{renew.}}, P^{f,\text{conv.}}, \lambda^{f,\text{DA}}, \lambda^{f,\text{bal.}}, \mathbf{x}^{f,\text{cal.}}\}$.

The level of relevance of input variables for predicting a target can be hardly anticipated. Hence, we train dedicated subnetworks, i.e., the variable selection layers, jointly with the model to filter out any irrelevant input. This process is showcased in Fig. 6.3 for the past observed inputs $\mathbf{x}_{t_0-l}^h$ at time step $t_0 - l$. First, each group within the inputs $\mathbf{x}_{t_0-l}^h$ is mapped into a d -dimensional vector, either linearly for the continuous variables or through entity embeddings for the calendar information [208]. The entity embeddings learn to map each calendar information to numerical features in a d -dimensional space. In contrast to the one-hot encoding methodology, this continuous representation identifies and

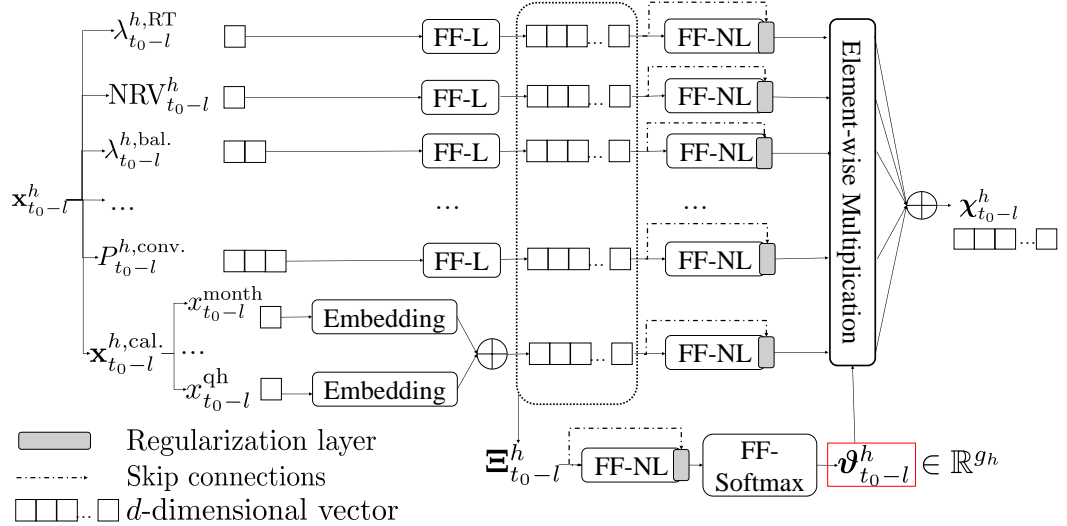


Figure 6.3.: Variable selection layer for the time step $t_0 - l$.

leverages similarities between time steps. Then, all the embedding vectors are averaged in a unique d -dimensional vector that condenses all the calendar information. Using a common representation space \mathbb{R}^d throughout the model enables skip connections, which facilitates the training (see Subsection 6.3.5).

The vector $\Xi_{t_0-l}^h \in \mathbb{R}^{g_h \cdot d}$ in Fig. 6.3 represents the concatenation of all the transformed inputs. Once non-linearly transformed, this vector is used as a basis to compute the feature importance variables $\vartheta_{t_0-l}^h$, framed in red in Fig. 6.3. They are obtained via a feed-forward network with a softmax function that outputs a vector of g_h -dimension. The softmax function ensures that the values of the output vector sum up to 1 and be positive. The final d -dimensional input $\chi_{t_0-l}^h$ for the time step $t_0 - l$ is then obtained by combining each transformed group of inputs, weighted by their corresponding value in $\vartheta_{t_0-l}^h$. Hence, the elements of the vector $\vartheta_{t_0-l}^h$ yield a probability distribution of the relative importance of each group in $\chi_{t_0-l}^h$, thereby providing interpretable outcomes (see Table 6.2 of Subsection 6.5.3).

6.3.3. Local Temporal Processing Layer

The input sequences χ^h, χ^f are then respectively processed by two distinct BLSTM networks, whose internal representations are exchanged at the forecast creation time t_0 . The BLSTM is composed of two LSTM networks that process the input sequence in both positive and negative time directions, which allows to capture both forward and backward local time dependencies. Without loss

of generality, the output of the BLSTM for the time step $t_0 - l$ is expressed as:

$$\mathbf{h}_{t_0-l}^{\text{forward}} = \mathcal{H}^{h,LSTM}(\mathbf{x}_{t_0-l}^h, \mathbf{h}_{t_0-l-1}^{\text{forward}}), \quad (6.4a)$$

$$\mathbf{h}_{t_0-l}^{\text{backward}} = \mathcal{H}^{h,LSTM}(\mathbf{x}_{t_0-l}^h, \mathbf{h}_{t_0-l+1}^{\text{backward}}) \quad (6.4b)$$

$$\mathbf{v}_{t_0-l}^h = \frac{\mathbf{h}_{t_0-l}^{\text{forward}} + \mathbf{h}_{t_0-l}^{\text{backward}}}{2} \quad (6.4c)$$

where \mathcal{H}^{LSTM} is the composite LSTM function [121] and $\{\mathbf{h}_t^{\text{forward}}, \mathbf{h}_t^{\text{backward}}\}$ are the internal states of the LSTMs. The output $\mathbf{v}_{t_0-l}^h$ is an average of both forward and backward internal states for keeping the same d -dimensional representation throughout the model.

Overall, the roles of the BLSTMs are to provide i) an appropriate information concerning the time ordering of the input sequence, and ii) awareness of the surrounding elements in the input sequence. Leveraging both time position and local context have proved to be key elements for computing the attention scores in the Transformer-specific attention layer [209].

6.3.4. Transformer-Specific Attention Layer

Attention mechanisms are computing layers that provide an abstract representation of an input sequence by dynamically weighting its different time steps. The process is showcased in Fig. (6.4) for the time step $t_0 + \tau$, where the sequence $\phi^{\{h,f\}} \in \mathbb{R}^{T \times d}$ (with $T = l_{\max} + \tau_{\max}$) is obtained from $\mathbf{v}^{\{h,f\}}$ using FF-NL layers. The sequence $\phi^{\{h,f\}}$ is linearly transformed in three different vectors, i.e., a query $\mathbf{Q}_{t_0+\tau} \in \mathbb{R}^d$, keys $\mathbf{K} \in \mathbb{R}^{T \times d}$ and values $\mathbf{V} \in \mathbb{R}^{T \times d}$, via FF-L layers. The abstract representation $\mathbf{A}_{t_0+\tau} \in \mathbb{R}^d$ is then obtained by weighting the values \mathbf{V} with attention scores $\alpha_{t_0+\tau} \in \mathbb{R}^T$, obtained by quantifying the level of matching between the query $\mathbf{Q}_{t_0+\tau}$ and the keys \mathbf{K} :

$$\mathbf{A}_{t_0+\tau} = a(\mathbf{Q}_{t_0+\tau}, \mathbf{K})\mathbf{V} \quad (6.5)$$

where $a(\cdot)$ is the matching function.

Following [203], we use the scaled dot-product attention as the matching function $a(\cdot)$:

$$a(\mathbf{Q}_{t_0+\tau}, \mathbf{K}) = \text{softmax}\left(\frac{\mathbf{Q}_{t_0+\tau}\mathbf{K}}{\sqrt{d}}\right) \quad (6.6)$$

The dot-product yields the similarity of vector $\mathbf{Q}_{t_0+\tau}$ with regard to the keys \mathbf{K} . Higher values of the dot-product correspond to higher rele-

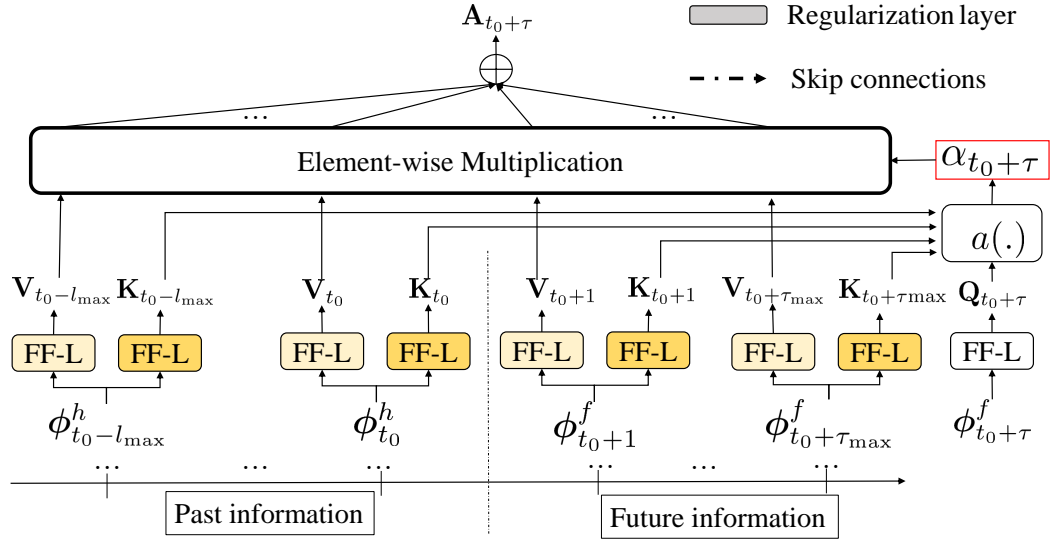


Figure 6.4.: The transformer-specific attention layer for the time step $t_0 + \tau$.

vance between the given key and the proposed query. The scaling factor $\sqrt{d_k}$ is introduced to reduce the magnitude of the dot-product. Then, the softmax function renders the attention scores $\alpha_{t_0+\tau}$ as a probability distribution over all keys \mathbf{K} with regards to $\mathbf{Q}_{t_0+\tau}$. The magnitude of the attention scores $\alpha_{t_0+\tau}$ provides direct insights on the contributions of each element of the input sequence $\phi^{\{h,f\}}$ to predict the real-time price at $t_0 + \tau$.

The attention mechanism provides two key benefits: i) the model is able to directly access to the most salient contextual information for each time step of the prediction horizon, and ii) it allows to learn regime-specific temporal dynamics by using distinct attention score patterns for each regime. These two benefits are respectively showcased in Fig. 6.10 and Fig. 6.11 of Subsection 6.5.3.

6.3.5. Regularization layer

To control the complexity and facilitate the training of the model, we use a regularization layer whenever a non-linear transformation is performed. This is illustrated in Fig. 6.5 for the time step $t_0 - l$ when using the BLSTM. The regularization layer is composed of three modules, i.e., a layer normalization [205], a gating mechanism [186] and a skip connection [210].

Layer normalization aims at speeding-up the training of the model by

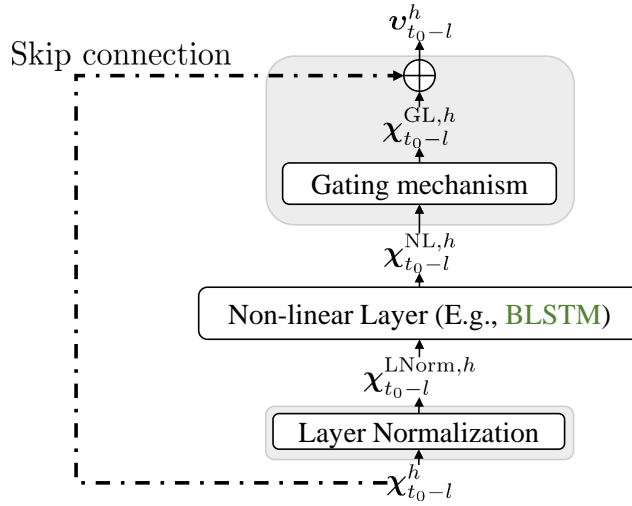


Figure 6.5.: Illustration of the regularization layer applied to the BLSTM for the time step $t_0 - l$.

normalizing the inputs at each layer. For this example, the output vector of the layer normalization is computed as $\mathbf{x}_{t_0-l}^{\text{LNorm},h} = \gamma \frac{\mathbf{x}_{t_0-l}^h - \mu}{\sigma} + \beta$, in which μ , σ are the mean and standard deviation of the elements in $\mathbf{x}_{t_0-l}^h$, and γ , β are the gain and bias parameters to be trained, respectively.

The gating mechanism allows the model to control the magnitude of the non-linear transformation of the previous layer. The gating mechanism, which takes as input $\mathbf{x}_{t_0-l}^{\text{NL},h}$, yields:

$$\mathbf{x}_{t_0-l}^{\text{GL},h} = f^\sigma(W^4 \mathbf{x}_{t_0-l}^{\text{NL},h} + b^4) \odot (W^5 \mathbf{x}_{t_0-l}^{\text{NL},h} + b^5) \quad (6.7)$$

where f^σ is the sigmoid function, $W^{\{4,5\}} \in \mathbb{R}^{d \times d}$, $b^{\{4,5\}} \in \mathbb{R}^d$ are weights and biases, \odot is the element-wise Hadamard product and d is the dimension of the model. If necessary, the gating mechanism could suppress the non-linear transformation by outputting values all close to 0.

Skip connections allow the model to learn residual functions, which have been proved to be easier to optimize in deeper architecture [210]. The skip connection simply performs an identity mapping, which is added to the output of the gating mechanism (neither extra parameters nor computational complexity is added). Overall, the gain of performance of using the regularization layers is analyzed in Fig. 6.9 of Subsection 6.5.2.

6.3.6. Output layer

The simultaneous prediction of the q -quantiles $\{\hat{\lambda}^{\text{RT}}, \forall q \in Q\}$, with Q the set of quantiles to predict, are achieved by a FF-L layer at each time step. To produce these quantiles, the model is trained using the smooth approximation of the pinball loss [143], where the Huber norm $H(\cdot)$ is introduced for differentiability issues at the origin. The loss function is computed as:

$$E_{\tau}^Q = \sum_{q \in Q} \begin{cases} q \cdot H(\lambda_{t_0+\tau}^{\text{RT}}, \hat{\lambda}_{t_0+\tau,q}^{\text{RT}}) & \hat{\lambda}_{t_0+\tau,q}^{\text{RT}} < \lambda_{t_0+\tau}^{\text{RT}} \\ (1-q) \cdot H(\lambda_{t_0+\tau}^{\text{RT}}, \hat{\lambda}_{t_0+\tau,q}^{\text{RT}}) & \hat{\lambda}_{t_0+\tau,q}^{\text{RT}} \geq \lambda_{t_0+\tau}^{\text{RT}} \end{cases} \quad (6.8)$$

As in Chapter 3, naive rearrangement of the predicted q -quantiles is conducted in ex-post, i.e., we sort in ascending order the q -quantiles at each time step of the prediction horizon after they are predicted [129]. This procedure is also performed for the benchmark.

The learning procedure is carried out with the Adam optimizer, with $\beta_1 = 0.9$, $\beta_2 = 0.98$ and $\epsilon = 10^{-9}$ [203]. The learning rate δ varies over the number n of mini-batches (a mini-batch is here composed of 96 sequences, and $n = 920$ is equals to one pass over the entire training set), according to the formula:

$$f^{\text{LR}}(\cdot) = \frac{\delta}{\sqrt{d}} \min \left(\frac{1}{\sqrt{n}}, \frac{n}{n_{\text{warmup}}^{1.5}} \right) \quad (6.9)$$

where d is the dimension of the model, while $\delta = 0.001$ and $n_{\text{warmup}} = 4000$ are hyperparameters that determine the highest learning rate achieved and the number of steps to reach it, respectively.

This function-based learning rate is shown in Fig. 6.6, which firstly shows a linear warmup followed by an inverse square root decay. The warmup phase allows the model to gradually learn weight parameters without triggering gradient explosion/vanishing issues. Then, higher learning rate values help to regularize the model during training by escaping, e.g., saddle points (but too large values may cause the training to diverge). Finally, the decay policy allows converging towards a minimum [205].

6.4. Benchmark

The proposed model is compared with the same time series forecasting models presented in Chapter 3, at the exception that only the two best neural networks are included. They are recalled below for the reader convenience:




Figure 6.6.: The function-based learning rate scheduler $f^{\text{LR}}(\cdot)$.

- The step-wise averaging model (Step-Avg).
- The probabilistic generalization of persistence (Persistence).
- The Auto-Regressive Moving Average (ARMA) model.
- The quantile regression forest (QRF).
- The gradient boosting regression tree (QGBRT).
- The traditional sequence-to-sequence model (Seq2Seq).
- The Bahdanau-based sequence-to-sequence model (Bahd-Seq2Seq).

The ARMA model is only fed with past imbalance price observations, while other machine learning models, i.e., QRF, QGBRT, Seq2Seq, and Bahd-Seq2Seq, have access to the same input data as the proposed transformer model. Similarly to Chapter 3, we conduct a hyperparameter optimization to identify the most suited model complexity of each forecaster. This is achieved through a random search, where the same number of iterations is used across all benchmarks.

Besides, we also perform an ablation study, in which we investigate the loss in performance of the proposed model (denoted by Ref) when removing important parts of its architecture.

- Ref-Att is the Ref model without the transformer-specific attention layer.

- Ref-VarSel is the Ref model where the variable selection networks are removed.
- Ref-Bidir is the Ref model where the sequential information fed to the attention mechanism is injected via sinusoidal functions of different frequencies [203] instead of the BLSTM networks.
- Ref-Reg is the Ref model without the regularization layers.

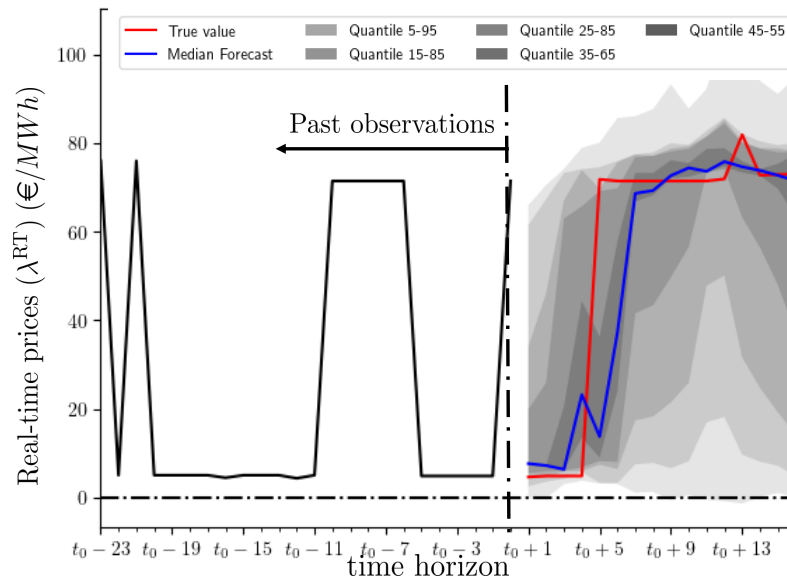
6.5. Case Studies

We conduct the case study on publicly available data obtained from the website of Elia **Elia web**, i.e., the Belgian Transmission System Operator, on an Intel® Core™ i7-3770 CPU @ 3.4 GHz with 16 Gb of RAM. The variable of interest is the Belgian imbalance price λ^{RT} . The twelve forecasting models are implemented using the scikit-learn package, statsmodels package, and TensorFlow package in Python 3.6. The data spans from 2016-1-1 to 2019-12-31, for a total of four years of data. Specifically, the first three years of data are used to train and validate the models with a ratio of 85%-15%, while the fourth year is used for testing. Each quarter-hourly step of the database is used as a forecast creation time t_0 . A prediction horizon of 4 hours is selected, which corresponds to $\tau_{\text{max}} = 16$ time steps, and we compute the 5th, 15th, 25th, 35th, 45th, 50th, 55th, 65th, 75th, 85th, 95th percentiles of the target distribution (i.e., $|Q| = 11$) for each of these time periods. The past conditioning range l_{max} is set to 24.

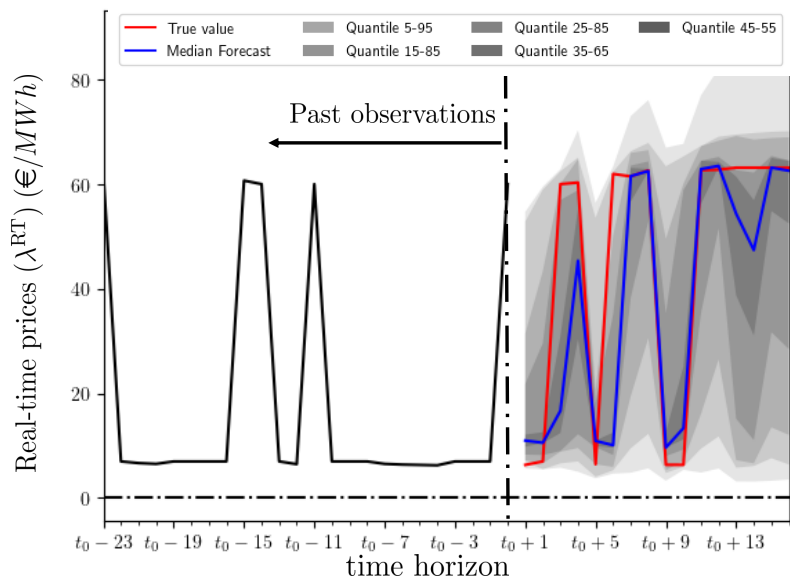
6.5.1. Forecast Evaluation

Fig. 6.7 illustrates the probabilistic forecasts obtained using the proposed model for the 14th April 2019 at 06H00 and the 14th September at 16H00. It can be observed that the real-time price signal is properly embedded by the prediction intervals. Interestingly, larger prediction intervals encompass both price regimes, while narrower intervals, e.g., $\{\hat{\lambda}_{t_0+\tau,0.35}^{\text{RT}}, \hat{\lambda}_{t_0+\tau,0.65}^{\text{RT}}\}$, attempt to focus on the accurate future price regime.

Table 6.1 provides the average continuous ranked probability score (CRPS) scores (defined in Section 3.8 of Chapter 3) of the different models for the entire forecasting horizon, where the best individual scores are denoted in bold figures. We observe that the proposed model (Ref) provides the lowest CRPS scores over the entire prediction horizon, while the other Machine Learning (ML) methods, i.e., Bahd-Seq2Seq, Seq2Seq, RF and GBDT, are the second-best models. One reason explaining the gap between ML methods



(a) 14th April 2019 at 06:00



(b) 14th September 2019 at 16:00

Figure 6.7.: Multi-horizon probabilistic forecasts of λ^{RT} on the 14th April 2019 at 06:00 (Fig. 6.7a) and on the 14th September 2019 at 16:00 (Fig. 6.7b).

Table 6.1.: Evolution of the CRPS score over the entire forecasting horizon for all the models, where Tot. is the aggregation of the CRPS scores.

Models	Tot.	$t_0 + 1$	$t_0 + 2$	$t_0 + 3$	$t_0 + 4$	$t_0 + 5$	$t_0 + 6$	$t_0 + 7$
Step-Avg	18.2	18.2	18.2	18.2	18.2	18.2	18.2	18.2
Persistence	23.2	18.47	20.93	21.61	21.4	22.59	23.43	23.74
ARMA	20.92	17.17	19.22	19.86	20.24	20.84	21.18	21.35
RF	18.67	15.69	17.27	17.75	18.02	18.67	18.92	19.01
GBDT	16.82	13.07	15.18	15.86	16.37	16.89	17.24	17.32
Seq2Seq	16.81	15.39	16.36	16.71	16.9	16.98	17.05	17.06
Bahd-Seq2Seq	16.34	14.7	15.88	16.14	16.25	16.42	16.51	16.53
Ref	15.6	12.88	14.5	15.2	15.54	15.79	15.98	16.05

	$t_0 + 8$	$t_0 + 9$	$t_0 + 10$	$t_0 + 11$	$t_0 + 12$	$t_0 + 13$	$t_0 + 16$
Step-Avg	18.2	18.2	18.2	18.2	18.2	18.2	18.2
Persistence	23.74	24.24	24.43	24.41	24.34	24.34	24.72
ARMA	21.44	21.59	21.62	21.63	21.66	21.67	21.79
RF	19.08	19.16	19.17	19.23	19.27	19.35	19.43
GBDT	17.45	17.5	17.4	17.36	17.47	17.57	17.47
Seq2Seq	17	16.96	16.94	16.91	16.94	16.93	16.98
Bahd-Seq2Seq	16.52	16.49	16.48	16.46	16.52	16.54	16.74
Ref	16.04	16.02	15.95	15.97	15.95	15.9	15.98

and the other naive and econometrics methods is that only ML methods fully leverage all the available input information. This tends to indicate that including exogenous variables in the forecasting models has a positive impact on accuracy. Interestingly, the naive Step-Avg model achieves an overall better performance than the autoregressive models, i.e., the Persistence and ARMA models. It is aligned with the observations in Chapter 3, where naive forecasts can be hard to beat for real-time market variables. The Persistence model has the worst performance within the benchmark. By simply propagating the most recent past realization, the model does not have the ability to infer the most likely future price-regime of the real-time prices. Even if it includes a larger look-back window of past realizations, the ARMA model is unable to perform better than the naive Step-Avg model. Overall, it can also be observed that the CRPS scores for all models (except for the Persistence) saturate when the prediction horizon is longer than one hour and a half, i.e., for t_{0+6} .

Concerning the ML models, we can see that the neural models, i.e., the Ref, Bahd-Seq2Seq, and Seq2Seq, outperform consistently the tree-based ensemble methods (RF and GBDT) over the last time steps $\{t_0 + 6, \dots, t_0 + 16\}$. This can be explained by the fact that different models are defined independently for each time step $t_0 + \tau$ for tree-based ensemble methods, whereas the

neural models are composed of a single model that is iterated over the entire forecasting horizon. By sharing the parameters in their output layer, the neural models are able to provide superior performance for longer horizon forecasts. However, it can be observed that the tree-based ensemble methods remain very competitive over the first five time steps and that the GBDT model performs even better than the Bahd-Seq2Seq. We also observe that RF performs worse than GBDT, which can be explained by the fact that RF gives an estimate of the conditional distribution from which quantiles are extracted, whereas the quantiles of GBDT are directly computed through the minimization of the quantile loss. Finally, both attention-based neural models, i.e., Ref and Bahd-Seq2Seq, provide the best averaged scores, which suggests that adding interpretable components within their architecture do not hinder their prediction performance.

To complement these results, we differentiate the performance of the models over the prediction intervals $\{\hat{\lambda}_{t_0+\tau,0.05}^{\text{RT}}, \hat{\lambda}_{t_0+\tau,0.95}^{\text{RT}}\}$, $\{\hat{\lambda}_{t_0+\tau,0.25}^{\text{RT}}, \hat{\lambda}_{t_0+\tau,0.75}^{\text{RT}}\}$, $\{\hat{\lambda}_{t_0+\tau,0.45}^{\text{RT}}, \hat{\lambda}_{t_0+\tau,0.55}^{\text{RT}}\}$ using the Winkler score, which respectively correspond to $\beta = \{0.1, 0.5, 0.9\}$. The results are showcased in Fig. 6.8 for the entire prediction horizon. For a larger interval at $\beta = 0.1$, the metrics of the ML models are very close to each other, and are significantly below the other models' metrics, which indicate that the ML models provide a sharper interval $\{\hat{\lambda}_{t_0+\tau,0.05}^{\text{RT}}, \hat{\lambda}_{t_0+\tau,0.95}^{\text{RT}}\}$. For such a large interval, the models provide predictions encompassing both the low- and high-price regimes (in a narrower fashion for the ML models), but none of the models are able to differentiate the price regime. Fig. 6.8b shows the Winkler score at $\beta = 0.5$. In this case, the metrics of Ref and GBDT are practically equal, while the Bahd-Seq2Seq and Seq2Seq perform worse over the first time steps. We observe that the performance of the RF model starts to deteriorate for narrower prediction intervals. Concerning the narrowest interval at $\beta = 0.9$ (Fig. 6.8c), the metrics of the models are more stratified. The Ref model clearly outperforms all other models, which is highly valuable since the 45th and 55th quantiles provide direct information on the price-regime of the real-time prices. This tends to demonstrate that the proposed model is able to better detect the likely future regime of real-time prices.

6.5.2. Ablation Analysis

Fig. 6.9 shows the average CRPS scores over the entire forecasting horizon resulting from the ablation analysis. First, we can see the importance of the regularization layers in training our proposed model to facilitate the convergence of the model. Indeed, the Ref-Reg model achieves the worst performance. The

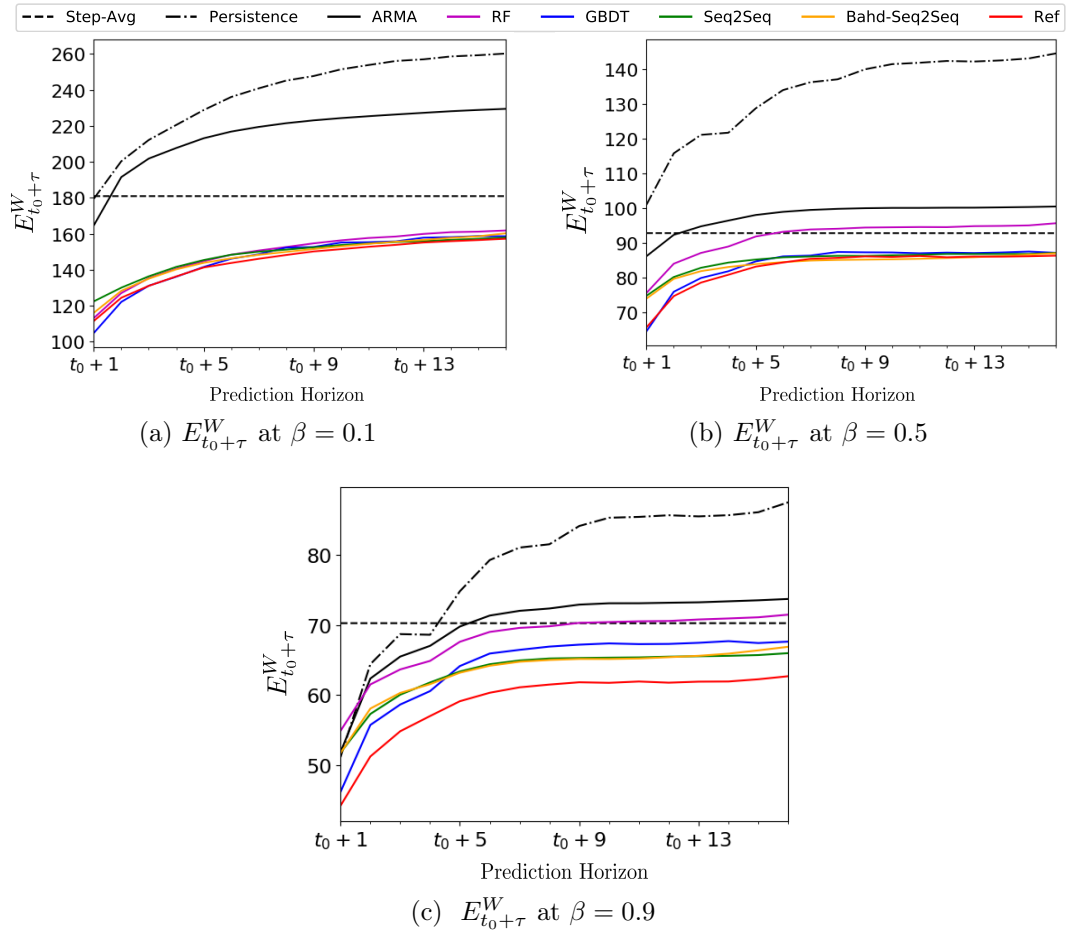


Figure 6.8.: Winkler score $E_{t_0+\tau}^W$ of all models for $\beta = \{0.1, 0.5, 0.9\}$ over the entire prediction horizon.

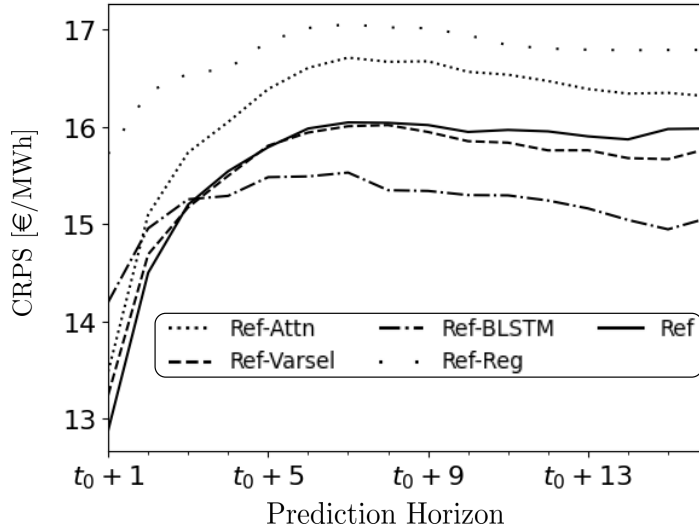


Figure 6.9.: CRPS scores for the ablation analysis over the entire prediction horizon.

non-attentional model, denoted Ref-Attn, is the second-worst model in terms of accuracy, which highlights the benefits of the alignment procedure that provides direct connections with relevant time steps of the surrounding horizon. This observation tends to reflect the importance of the attention mechanism to capture different temporal patterns for differentiating the different regimes of price. Interestingly, we see that injecting the sequential information with sinusoidal functions of different frequencies (instead of the BLSTM) also worsens the results during the first quarter hours, but provides better forecasts for the remaining prediction horizon. Finally, the metrics of the models Ref-VarSel and Ref are comparable over the entire forecasting horizon. This suggests that the main interest of adding the variable selection layers consists in providing more interpretable outcomes.

6.5.3. Interpretability

The interpretable outcomes of the variable selection layers are shown in Table 6.2. Practically, $\{\bar{v}^h, \bar{v}^f\}$ are selection variables that are aggregated for each feature across the entire test set. Results show that the proposed model extracts only a subset of key inputs (highlighted in bold) that intuitively play a significant role in the predictions. Regarding the past available information, the lagged values of the real-time prices are critical as expected. In addition, the net activated regulation volume and the calendar information also emerge

Table 6.2.: Averaged representation of variable selection weights over both past (left column) and future (right column) data.

	$\bar{\vartheta}^h$	$\bar{\vartheta}^f$
λ^{RT}	0.27	/
NRV	0.26	/
$\lambda^{\text{bal.}}$	0.05	/
ϕ	0.03	/
$L^{\{h,f\}}$	0.01	0.03
$P^{\{h,f\},\text{renew.}}$	0.11	0.04
$P^{\{h,f\},\text{conv.}}$	0.05	< 0.01
$x^{\text{cal.}}$	0.22	0.02
λ^{DA}	/	0.86
$\lambda^{\text{f,bal.}}$	/	0.05

as important drivers for the model. Remarkably, the lagged values of renewable generation bring an additional explanation power to the model. For the known inputs in $\mathbf{x}_{t_0+1}^f$, the most dominant driver is the day-ahead electricity price. However, in our experimental set-up, the merit order proxies of operational balancing reserves provided by the TSO, i.e., $\lambda^{\text{f,bal.}}$, play only a minor role in the proposed model.

Next, we analyze persistent temporal patterns, which are often key to understanding the time-dependent relationships between inputs-outputs. To do so, we average the attention weights over the entire test set, which produces the averaged attention patterns $\bar{\alpha}_{t_0+\tau}$ depicted in Fig 6.10. In this plot, each contour line perpendicular to the ‘conditioning range’ axis represents the intensity of the model’s temporal attention for each time step of the prediction horizon. Over the whole prediction horizon, it can be seen that the model is mostly focused on the time steps between $t_0 - 7$ and $t_0 + 16$. Such outcomes can be expected since the real-time price is a signal which includes quick and abrupt changes. It should be noted that Table 6.2 and Fig. 6.10 are also useful for model designers, which can use these interpretable outcomes for improving the model. Indeed, based on Table 6.2, some irrelevant covariates could be removed from the model, e.g., the variable $P^{\text{f,conv.}}$. Furthermore, the outcomes of Fig. 6.10 can be used for fine-tuning the length of the past conditioning




Figure 6.10.: Averaged temporal attention of the model over the entire prediction horizon for both past and future conditioning ranges

range, i.e., a 2-hour look-back window seems already sufficient for our case study.

Finally, we conduct a case-specific interpretable analysis in Fig. 6.11 for the prediction time steps $\{t_0 + 1, t_0 + 5, t_0 + 13\}$ of the probabilistic forecasts on 14th April 2019 at 06H00. Recalling Fig. 6.7a, the forecaster predicts a low-price regime at $t_0 + 1$, then, it introduces a regime switch at $t_0 + 5$ in order, finally, to output a high-price regime distribution at $t_0 + 13$. Fig. 6.11 clearly demonstrates the dependency between the predicted regimes and the different temporal importance patterns of the most dominant drivers. Concretely, the temporal importance of each dominant driver is computed as $\varrho_{t_0+j}^i = \vartheta_{t_0+j}^i \cdot \alpha_{t_0+j}$, $\forall i \in \{\lambda^{h,RT}, NRV^h, P^{h,renew}, x^{h,cal.}, \lambda^{f,DA}\}$, $\forall j \in \{-l_{max}, \dots, \tau_{max}\}$. It can be observed that the proposed model tends to rely on the day-ahead prices for predicting a low imbalance price, while it focuses on past information for predicting a high-price regime. These observations are further corroborated by the importance spikes occurring on the conditioning steps $\{t_0, t_0 - 7, t_0 - 21\}$ in Fig. 6.11b and 6.11c, which correspond to past time steps characterized by a high-price regime (see Fig. 6.7a).

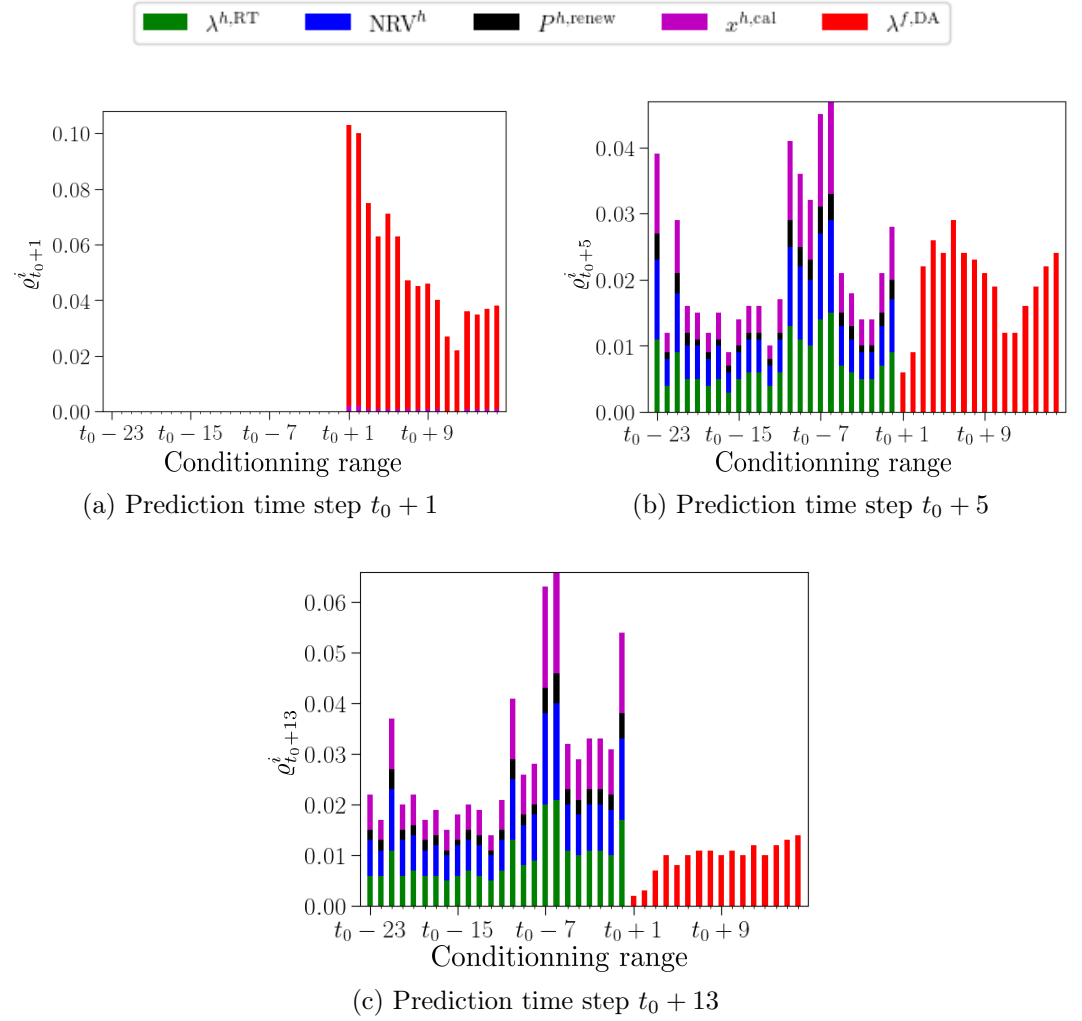


Figure 6.11.: The attention-weighted importances of the most dominant drivers over both past and future conditioning ranges, i.e., $\varrho_{t_0+j}^i = \vartheta_{t_0+j}^i \cdot \alpha_{t_0+j}$, $\forall i \in \{\lambda^{\text{RT}}, \text{NRV}, P^{\text{h},\text{renew}}, x^{\text{cal}}, \lambda^{\text{DA}}\}$, $\forall j \in \{-l_{\text{max}}, \dots, \tau_{\text{max}}\}$, when performing the probabilistic forecasts on the 14th April 2019 at 06H00 for the prediction time steps $\{t_0 + 1, t_0 + 5, t_0 + 13\}$ (which respectively corresponds to Fig. 6.11a, 6.11b, and 6.11c).

6.6. Conclusion

This Chapter proposes a novel Transformer-based model for interpretable, high-performance multi-horizon probabilistic forecasting of real-time electricity prices. Such prices are important market signals for market players aiming at reducing their imbalance costs or maximizing balancing actions. However, their predictions are highly complex since the prices are characterized by regime switching behavior and spikes.

In this context, we leverage recent advances in deep neural networks to provide a new, well-suited approach for predicting real-time electricity prices. In a detailed case study, we illustrate that the proposed model is able to outperform state-of-the-art forecasting methods, with a respective decrease of 4.5% in the CRPS metrics compared with the first benchmark methods, i.e., Bahd-Seq2Seq. Results suggest that performance and accuracy in deep learning-based time series forecasting model can be jointly improved. In particular, the ablation analysis show that enriching the model with attributes dedicated to yield more interpretable outcomes do not come at the expense of performance in our case study. In addition, a global analysis of the interpretable outcomes allow highlighting both the most important features, i.e., the set $\{\lambda^{h,RT}, NRV^h, P^{h,renew}, x^{h,cal}, \lambda^{f,DA}\}$, and the features' temporal window of the proposed model (from $t_0 - 7$ to $t_0 + 16$). Finally, a case-specific interpretable analysis demonstrates the ability of the proposed model to capture different temporal attention patterns of each input according to the price-regime predicted. Such an interpretable analysis, based on a purely data-driven approach, could be replicated in another market framework in order to investigate different real-time market dynamics.

Chapter Publications

- **J. Bottieau**, Y. Wang, Z. De Grève, F. Vallée, and J-F. Toubéau, "Transformer Model for Interpretable Probabilistic Forecasting of Real-Time Electricity Prices," currently in second round of review in IEEE Trans. Power Syst.
- J-F. Toubéau, **J. Bottieau**, Y. Wang, and F. Vallée, "Interpretable Probabilistic Forecasting of Imbalances in Renewable-Dominated Electricity Systems," in IEEE Trans. Sust. Energy. Early access: <https://ieeexplore.ieee.org/abstract/document/9464660>.

CHAPTER 7.

Conclusions and Perspectives

This final chapter summarizes the main contributions and findings of this report, and gives suggestions for further research.

7.1. Conclusions

In this dissertation, novel forecast-driven strategies were developed for fostering the provision of real-time balancing services in European electricity markets. Practically, these strategies are studied using an integrated approach, where the entire value chain, i.e., from forecasting to the decision-making processes, is modeled for optimizing close-to-real-time the imbalance position of a Balance Responsible Party (BRP) within the single price imbalance settlement. Firstly, probabilistic neural models of gradual complexity have been proposed for forecasting the system imbalance, aiming at better capturing the temporal dependencies between past observed and future known inputs. In this line, a sequence-to-sequence Long Short Term Memory neural architecture augmented with attention mechanisms has showcased high-quality predictions compared to other competitive prediction models (e.g., tree-based ensemble methods, feed-forward neural networks or sequence-to-sequence neural models). Then, for exploiting at best the predicted uncertain market signal, stochastic decision-support tools, i.e., stochastic programming and robust optimization model, have been formulated for guiding the imbalance strategy of the BRP. Results suggest that i) owning high-quality probabilistic forecasts of the system imbalance, ii) including the impact of the BRP on the balancing energy market-clearing process, and iii) modeling the uncertainty in the decision-support tool, improve the economic gains of the BRP. Indeed, the single price imbalance settlement is not a market *per se*, but rather a mechanism aiming at supporting the real-time balancing of the system, and, consequently, may lead to financial penalties for the unbalanced BRPs. In

that direction, for a better provision of real-time balancing services, a new data-driven approach was designed for continuously adjusting the risk policy of the BRP decision-support tool. The resulting automatic risk-adjusted policy shows promising economic benefits in comparison with a static (determined once and for all) risk policy. In all case studies based on real-world market data of the Belgian power system, the implementation of these forecast-driven strategies corroborates the key goal of the single price imbalance settlement, by showing that the BRP can increase its operating profit by optimizing its imbalance position, while reducing the total net system imbalance. Finally, research efforts are devoted to improve the designer and/or user confidence in probabilistic forecasts based on deep neural models, by adding interpretability pathways within their architecture.

Recalling **Chapter 1**, this PhD thesis sets the following four research objectives:

1. **Attaining high-quality probabilistic forecasts.** High-quality probabilistic forecasts of the total net system imbalance are a mandatory condition for a well-informed out-of-balance position of the market actor. We therefore investigate novel deep learning-based time series forecasting tools for generating improved predictions.
2. **Developing tailored risk-aware stochastic decision-support tool.** Based on the predicted probabilistic information, novel mathematical formulations are developed for accommodating i) the small energy volumes exchanged in real time, and ii) the high uncertainty and risk associated with these real-time opportunities.
3. **Adjusting continuously the risk policy of a market actor.** In this objective, machine learning techniques are leveraged for adjusting, at each decision step, the risk policy of the Balancing Responsible Party based on the current state of its expected market outcomes.
4. **Adding interpretability in probabilistic forecasts.** This objective aims at incorporating notions of interpretability (i.e., the identification of the most important input features over time) in deep learning-based time series forecasters for improving the user confidence in their outcomes.

Chapter 2 proposes a global presentation of the current status of the European electricity markets, while presenting our market application, i.e., the real-time provision of balancing services. Then, each of the following Chapters follows a one-to-one mapping in relation with each research objective. Short highlights of each Chapter are given below, providing an overview of the proposed research contributions.

Chapter 3 - High-quality Probabilistic Forecasts

Chapter 3 generates multi-horizon probabilistic forecasting of the system imbalances based on neural networks of gradual complexity. The system imbalance is a highly variable and uncertain signal, which directly results from the real-time operating conditions of the system. The anticipation of this market signal is essential for Balancing Responsible Parties (BRPs) that aim at providing real-time balancing services. Indeed, the direction of the system imbalance is in direct relation with the imbalance price regime (either low via the Marginal Decremental balancing energy Price (MDP) or high via the Marginal Incremental balancing energy Price (MIP)). The following neural architectures are assessed and compared: i) shallow (i.e., 1-layer) and stacked feed-forward neural networks, ii) Long Short Memory recurrent neural networks and their bidirectional counterpart, and iii) the sequence-to-sequence with and without attention mechanisms. All these neural architectures have been detailed, and are sequentially suggested for better capturing the specificities pertaining to time series forecasting (whose inputs can be composed of past observed and future known inputs). The outcomes show that advancements in terms of neural architecture are accompanied with an increase of forecasting performance. Practically, the sequence-to-sequence Long Short Term Memory neural networks augmented with attention mechanisms shows the highest accuracy compared to other benchmark outcomes. This suggests that tailoring the neural architectures to the temporal specificities of time series forecasting allows to generate more accurate (tightened) quantiles.

Chapter 4 - Risk-aware Stochastic Decision-Support Tool.

For fostering the provision of real-time balancing services, Chapter 4 proposes a risk-aware stochastic decision-support tool, which i) is based on a bi-level structure, that allows mathematically capturing the interaction of the BRP with the real-time system balancing, ii) includes the system imbalance uncertainty via robust-based or stochastic programming optimization formulations, and iii) proposes to manage the financial risk in both optimization formulations, i.e., by adjusting the uncertainty set for robust-based optimization, and by including the conditional Value-at-Risk (CVaR) within the objective function of the stochastic programming. Combining these three mathematical aspects allows responding to three challenges pertaining to our market application: i) the committed energy volumes in real-time are inherently small, ii) the system imbalance is highly uncertain, and iii) real-time balancing services aggravating the system imbalance incur financial losses. Based on extensive case studies using real-life market data from the Belgian power systems, outcomes suggest that i) the bi-level structure is efficient for hedging against the inherent

small volume of system imbalances and the associated imbalance price regime switching effect, ii) gains of accuracy in the probabilistic predictions allow achieving better decisions, and thus, better ex-post economic profits, and iii) a financial risk management (for conservative risk policies) achieves less economic losses when the market conditions are unsure, but at the expense of opportunities costs in favorable situations. This trade-off is more pronounced for the robust-based optimization, in comparison with the CVaR stochastic programming, whose ex-post economic profits decrease gradually with more conservative policies.

Chapter 5 - Automatic Risk Adjustment.

Chapter 5 uses Machine Learning techniques (i.e., linear regression, feed-forward neural networks, tree-based ensemble methods and k -nearest neighbours) to dynamically and preemptively adjust the risk policy of a Balancing Responsible Party (BRP) when providing real-time balancing services. This automatic risk-adjusted approach differs from the commonly proposed risk policies in the literature, which are time-invariant as based on a single computationally-intensive out-of-sample analysis. The proposed risk-adjusted approach is well-adapted for our market application as the real-time system operating conditions are changing between each decision step (e.g., via the ramping trajectories due to the hourly time resolution of the day-ahead market). This approach is tested on two risk-aversion parameters, i.e., the confidence level of the conditional Value-at-Risk (CVaR) and the budget of uncertainty, respectively considering both the CVaR-based stochastic programming and robust optimization frameworks. Both automatic risk-adjusted decision-support tools are then assessed and tested on extensive numerical analyses, using real-world market data from the Belgian power system over one year. Overall, the outcomes highlight the economic potential of adopting a time-dependent risk policy for both robust-based and CVaR-based stochastic optimizations in comparison with traditional time-invariant risk policies. More particularly, regarding robust-based formulation, a dynamic support for defining the uncertainty set is a key element for fostering its ex-post profits, with an increase of 12.3% (when guided by the k -NN technique) compared to the optimal time-invariant risk policy. In the same vein, the CVaR stochastic program adjusted by the k -NN model has presented a moderate increase of 4.8% in the ex-post profits compared with the optimal time-invariant risk policy.

Chapter 6 - Towards Interpretable Probabilistic Forecasts.

Chapter 6 falls within the research line of combining the predictive power of deep neural models with interpretable features applied to time series. In short,

this Chapter aims at outputting accurate predictions, while identifying the most important input features of the associated neural model and their interaction. Developing an interpretable model can be beneficial for i) the end-user, who will be better equipped for making a decision, and ii) the designer, who can check that the model does not exploit artifacts in the data. This Chapter focuses on probabilistic imbalance price forecasting (instead of the system imbalance), which allows highlighting distinct temporal patterns of the input signals depending on the predicted price regime. The proposed interpretable model is based on Transformer, which tends to become the novel state-of-the-art neural model in various applications such as, e.g., natural language processing. By relying on attention mechanisms solely computed via feed-forward neural networks, the proposed model is designed to capture distinct temporal patterns of the input signal depending on the predicted price regime. In addition, the model is augmented with subnetworks able to provide direct insights on the relative importance of each individual input feature. A detailed case study based on Belgian data over 2016-2019 illustrates that the proposed model is able to outperform state-of-the-art forecasting methods, with a respective decrease of 4.5% in the continuous ranked probability score (CRPS) metric compared with the first benchmark method, i.e., the sequence-to-sequence Long Short Term Memory neural network augmented with attention mechanisms. In addition, a global analysis of the interpretable outcomes allow highlighting both the most important features and the features' temporal window of the proposed model. Finally, a case-specific interpretable analysis demonstrates the ability of the proposed model to capture different temporal attention patterns of each input according to the price-regime predicted.

7.2. Perspectives

Based on the current limitations of the proposed research work, this Section suggests several perspectives for future research. Firstly, three axes of progression are proposed concerning the decision-support tool providing the real-time balancing services. Then, different research gaps and perspectives are presented in regards with energy forecasting based on neural models.

On the Provision of Real-Time Balancing Services

The following three gaps and proposals are identified concerning the provision of real-time balancing services:

- **Perspective 1 - on the temporality of the decision-support tool.**
The proposed decision-support tool only considers a single look-ahead

time step in its formulation, thereby neglecting the fact that here-and-now decisions are affecting potential future rewards, and, consequently, may jeopardize future opportunities. Hence, a first natural extension of this formulation consists in enlarging the look-ahead horizon of the decision-support tool, which could provide valuable insights regarding temporal arbitraging opportunities. Indeed, by having a longer view on the market dynamics, the decision-support tool might anticipate future high reward - low risk opportunities, and timely secure larger margins of flexibility for those time steps. In complement, the current decision-support tool could also be integrated in a more global market optimization framework. For instance, the provision of real-time balancing services could be also combined with the continuous intraday market opportunities, which will permit the flexible Balance Responsible Party to diversify its streams of revenue with less riskier market opportunities.

- **Perspective 2 - on the techno-economic constraints of the BRP portfolio and intra-period operational scheduling** (e.g., with minute-wise granularity in decisions). In the proposed decision-support tool, the techno-economic constraints of the BRP portfolio are reduced to their simplest expression, with only upward and downward 120 MW power limits. In that direction, a more accurate modeling of the techno-economic constraints of the BRP portfolio could provide interesting insights on their practical impact on the provision of real-time balancing services. Indeed, for instance, a lithium-ion battery energy storage system, facing higher degradation cost in extreme state-of-charge regimes, may limit its provision strategy (e.g., by not going deep in its state-of-charge regime) for extending its service life. In addition, for increasing the total volume of its balancing services, a BRP may diversify its portfolio technologies, where i) fast response technologies (e.g., the lithium-ion battery) exploit their fast ramping capabilities at the initial response, which can then be relieved by ii) slower power technologies (e.g., thermal units), which are characterized by higher ramping constraints, for a longer activation duration. Such a strategy is associated with the development of a multi-time scale approach, where i) a very-short-time scale formulation (e.g., a 1-minute granularity) will model the bridging between fast and slower response technologies, and ii) a short-time scale (e.g., a 15-minute resolution) will take into account the profit opportunities that could be achieved in the following time steps.
- **Perspective 3 - on the automatic risk-adjusted approach.** There is always a gap between expected (in-sample) objective outcomes at different risk-aversion parameters and the corresponding out-of-sample objective

outcomes, i.e., their actual realizations, which has been highlighted in the case studies in Chapter 4. Such a gap can result from a mis-representation of the decision-support tool (e.g., inaccurate representation of the portfolio constraints or the market-clearing process) or a mis-modeling of the uncertainty. For closing this gap, Chapter 5 proposes a Machine Learning-based supervised framework trained on past trading sessions, which provides early estimates of the out-of-sample objective outcomes. This allows selecting the most adequate risk-aversion parameter based on the maximum value of the early estimates. Instead of a supervised framework, a Reinforcement Learning approach could be investigated for learning an optimal risk trading strategy. In this approach, fully data-driven decisions are given to the BRP, wherein the optimal trading policy is learned by directly interacting with a market simulator and/or the physical environment. This framework shows potential for reducing the two-step automatic risk-adjusted decision-support tool of Chapter 5 into a single decision-making problem, as the BRP reward would be directly computed based on the out-of-sample objective outcomes.

On Energy Forecasting

Beyond solely improving the model accuracy, research efforts are still required for improving the user interaction with deep learning-based solutions. In this line, three research perspectives are suggested for fostering their practical adoption in energy industries:

- **Perspective 1 - on interpretability in neural forecasters.** Although some notions of interpretability have been touched in Chapter 6, this research line in time series forecasting deserves more in-depth studies. For instance, an extensive benchmark study investigating both post-hoc and inherent interpretable methods could provide interesting insights on their usefulness for providing human interpretable outcomes in time series forecasting applications. Besides, interpretability in time series forecasting may suffer from low intelligibility, i.e., their interpretations are not easily apprehended by non-expert users. This makes them mainly beneficial for experts rather than end-users. For leveraging intelligibility, natural language interpretations based on, e.g., a template approach, could be researched for directly mapping causal interpretations of the prediction outcomes into a linguistic structure, where gaps are filled from the raw interpretation. A simplistic example commenting a wind power forecast would be: the wind production will be [state]: e.g., high in afternoon, because we expect [observation]: e.g., a high wind speed.
- **Perspective 2 - on adaptability in neural forecasters.** In this

dissertation, forecasting models employ (mini-)batch learning, which results in estimating the model parameters once and for all on a training dataset. The predictions models are then deployed online to infer predictions. However, with a changing climate and energy policy frameworks, new patterns may progressively appear in the data related to the energy sector, which can lead to a deterioration of the model accuracy over time. In this setting, one of the key challenges is to correctly balance the trade-off between plasticity (i.e., the ability of the model to adapt to new knowledge) and stability (i.e., the ability of the model to retain prior knowledge). An excessive plasticity can cause catastrophic forgetting of persistent patterns, while an extreme stability may delay the learning of new temporal dynamics. For preventing catastrophic forgetting of the neural models, which can lead to a deterioration of the model accuracy over time, novel (online) training solutions can be envisaged. Promising alternatives can be i) the elastic weight consolidation algorithm, which constrains the most contributing parameters to stay close to their old values via a quadratic penalty, and ii) an episodic memory-based approach, where previous knowledge is consolidated by enforcing at least an equal performance of the model on a representative subset of previous data samples during the update.

- **Perspective 3 - on robustness in neural forecasters.** Outputting probabilistic predictions instead of deterministic ones certainly increases the robustness of the neural models by quantifying the uncertainty distribution around the predicted variable. However, several empirical observations show that neural models can be vulnerable to small perturbations in input data, which rises severe security concerns when applied in the field. A promising research approach for mitigating such perturbations (that can be intentional or not) could be provided by using an adversarial training scheme, which formulates a min-max regression problem when training the neural forecaster. This two-step approach minimizes the neural model in the outer loop, based on firstly crafted worst-case perturbations that are optimized the inner loop for hurting the model performance.

Bibliography

- [1] V. Smil, “Fossil-fueled civilization,” in *Energy and Civilization: A History*, Cambridge, MA: The MIT Press, 2017, ch. 6, pp. 295–384.
- [2] J. Bolt and J. L. van Zanden, “Maddison style estimates of the evolution of the world economy. a new 2020 update,” *Maddison Project Database, version 2020*, 2020.
- [3] United Nations Climate Change Conference, “Paris climate change conference - november 2015,” *21st Conference of the Parties*, 2015.
- [4] European Commission, *EU energy in figures: Statistical pocketbook 2021*. Luxembourg: Publications Office of the European Union, 2021.
- [5] The Intergovernmental Panel on Climate Change, “Summary for policymakers,” in *Climate Change 2021: The Physical Science Basis*, Cambridge University Press: Press, 2021.
- [6] European Commission, *Communication from the Commission to the European Parliament, the Council, the European economic and social Committee and the Committee of the regions empty. 'Fit for 55': delivering the EU's 2030 Climate Target on the way to climate neutrality*. Luxembourg: Publications Office of the European Union, 2021.
- [7] K. Keramidas, A. D. Vazquez, M. Weitzel, T. Vandyck, M. Tamba, S. Tchung-Ming, A. Soria-Ramirez, J. Krause, R. V. Dingenen, Q. Chai, S. Fu, and X. Wen, “The role of electrification in low-carbon pathways, with a global and regional focus on eu and china,” in *Global Energy and Climate Outlook 2019: Electrification for the low-carbon transition*, Luxembourg: Publications Office of the European Union, 2020.
- [8] H. Hoschle, *Capacity Mechanisms in Future Electricity Markets*. PhD thesis KU Leuven, 2018.
- [9] K. Brunninx, *Improved modeling of unit commitment decisions under uncertainty*. PhD thesis KU Leuven, 2016.

Bibliography

- [10] G. Papaefthymiou and K. Dragoon, "Towards 100% renewable energy systems: Uncapping power system flexibility," *Energy Policy*, vol. 92, pp. 69–82, 2016.
- [11] S. Impram, S. Varbak Nese, and B. Oral, "Challenges of renewable energy penetration on power system flexibility: A survey," *Energy Strategy Rev.*, vol. 31, pp. 100–539, 2020.
- [12] Elia Group, *Adequacy and flexibility study for Belgium 2020 - 2030*. 2019.
- [13] L. Duchesne, *Machine Learning of Proxies for Power Systems Reliability Management in Operation Planning*. PhD thesis ULiège, 2021.
- [14] California Independent System Operator, *What the duck curve tells us about managing a green grid*. 2013. [Online]. Available: <http://large.stanford.edu/courses/2015/ph240/burnett2/docs/flexible.pdf> (visited on 10/12/2021).
- [15] K. Van den Bergh and E. Delarue, "Cycling of conventional power plants: Technical limits and actual costs," *Energy Convers. Manag.*, vol. 97, pp. 70–77, 2015.
- [16] K. Bruninx, Y. Dvorkin, E. Delarue, W. D'haeseleer, and D. S. Kirschen, "Valuing demand response controllability via chance constrained programming," *IEEE Trans. Sust. Energy*, vol. 9, no. 1, pp. 178–187, 2018.
- [17] Z. Liu, Q. Wu, S. Huang, L. Wang, M. Shahidehpour, and Y. Xue, "Optimal day-ahead charging scheduling of electric vehicles through an aggregative game model," *IEEE Transact. Smart Grid*, vol. 9, no. 5, pp. 5173–5184, 2018.
- [18] Y. Mou, A. Papavasiliou, and P. Chevalier, "A bi-level optimization formulation of priority service pricing," *IEEE Trans. Power Syst.*, vol. 35, no. 4, pp. 2493–2505, 2020.
- [19] M. Hupez, J.-F. Toubeau, Z. De Grève, and F. Vallée, "A new cooperative framework for a fair and cost-optimal allocation of resources within a low voltage electricity community," *IEEE Trans. Smart Grid*, vol. 12, no. 3, pp. 2201–2211, 2021.
- [20] T. Brijs, *Electricity storage participation and modeling in short-term electricity markets*. PhD thesis KU Leuven, 2017.
- [21] Engie, *Coo pumped storage power station*. 2021. [Online]. Available: <https://corporate.engie.be/fr/energy/hydraulique/centrale-daccumulation-par-pompage-de-coo> (visited on 10/12/2021).

- [22] ———, *Engie Drogenbos project, a Bridge Between Business and Research*. 2018. [Online]. Available: <https://ease-storage.eu/news/engie-drogenbos-project/> (visited on 10/12/2021).
- [23] T. Gomez, I. Herrero, P. Rodilla, R. Escobar, S. Lanza, I. Fuente, M. Llorens, and P. Junco, “European union electricity markets: Current practice and future view,” *IEEE Power and Energy Magazine*, vol. 17, pp. 20–31, 2019.
- [24] J. Martínez, P. C. Kjær, P. Rodriguez, and R. Teodorescu, “Design and analysis of a slope voltage control for a DFIG wind power plant,” *IEEE Trans. Energy Convers.*, vol. 27, no. 1, pp. 11–20, 2012.
- [25] S. A. Hosseini, J.-F. Toubreau, Z. De Grève, and F. Vallée, “An advanced day-ahead bidding strategy for wind power producers considering confidence level on the real-time reserve provision,” *Appl. Energy*, vol. 280, p. 115973, 2020.
- [26] M. Gibescu, W. Kling, B. Ummels, E. Pelgrum, and R. van Offeren, “Case study for the integration of 12 GW wind power in the dutch power system by 2020,” in *2009 CIGRE/IEEE PES Joint Symposium*, 2009.
- [27] S. Leva, A. Nespoli, S. Pretto, M. Mussetta, and E. Ogliari, “Pv plant power nowcasting: A real case comparative study with an open access dataset,” *IEEE Access*, vol. 8, pp. 194428–194440, 2020.
- [28] A. M. Jafari, H. Zareipour, A. Schellenberg, and N. Amjady, “The value of intra-day markets in power systems with high wind power penetration,” *IEEE Trans. Power Syst.*, vol. 29, no. 3, pp. 1121–1132, 2014.
- [29] H. Ding, Z. Hu, and Y. Song, “Rolling optimization of wind farm and energy storage system in electricity markets,” *IEEE Trans. Power Syst.*, vol. 30, no. 5, pp. 2676–2684, 2015.
- [30] D. S. Kirschen and G. Strbac, *Fundamentals of Power System Economics, 2nd Ed.* West Sussex, U.K.: Wiley, 2018.
- [31] A. Papavasiliou and Y. Smeers, “Remuneration of flexibility using operating reserve demand curves: A case study of belgium,” *Energy J.*, vol. 38, 2017.
- [32] P. Shinde and M. Amelin, “A literature review of intraday electricity markets and prices,” in *2019 IEEE Milan PowerTech*, 2019, pp. 1–6.
- [33] ENTSO-E, *Single Intraday Coupling (XBID) Information Package*. 2018. [Online]. Available: https://www.nemo-committee.eu/assets/files/SIDC_Information%20Package_April%202021-99076f6ed5001c4d47442ae5cccebf30.pdf (visited on 10/12/2021).

Bibliography

- [34] ———, *ENTSO-E Balancing Report*. 2020. [Online]. Available: https://eepublicdownloads.azureedge.net/clean-documents/Publications/Market%20Committee%20publications/ENTSO-E_Balancing_Report_2020.pdf (visited on 10/12/2021).
- [35] European Commission, “The electricity balancing guideline - 2017/2195 of 23 november 2017,” *Official Journal of the European Union*, 2017.
- [36] T. Brijs, C. De Jonghe, B. F. Hobbs, and R. Belmans, “Interactions between the design of short-term electricity markets in the CWE region and power system flexibility,” *Appl. Energy*, vol. 195, pp. 36–51, 2017.
- [37] A. Abdisalaam, I. Lampropoulos, J. Frunt, G. P. J. Verbong, and W. L. Kling, “Assessing the economic benefits of flexible residential load participation in the dutch day-ahead auction and balancing market,” in *EEM 2012 : 9th International Conference on the European Energy Market*, 2012, pp. 1–8.
- [38] J. Z. Riveros, R. Donceel, J. V. Engeland, and W. D’haeseleer, “A new approach for near real-time micro-chp management in the context of power system imbalances – a case study,” *Energy Convers. Manag.*, vol. 89, pp. 270–280, 2015.
- [39] C. Koch, “Intraday imbalance optimization: Incentives and impact of strategic intraday bidding in germany,” *SSRN Electronic Journal*, 2019.
- [40] P. Wang, H. Zareipour, and W. D. Rosehart, “Descriptive models for reserve and regulation prices in competitive electricity markets,” *IEEE Trans. Smart Grid*, vol. 5, no. 1, pp. 471–479, 2014.
- [41] J.-F. Toubreau, J. Bottieau, Y. Wang, and F. Vallee, “Interpretable probabilistic forecasting of imbalances in renewable-dominated electricity systems,” *IEEE Trans. Sust. Energy*, 2021, Early access: <https://ieeexplore.ieee.org/document/9464660>.
- [42] J. Bottieau, K. Bruninx, A. Sanjab, Z. De Grève, F. Vallée, and J.-F. Toubreau, “Automatic risk adjustment for profit maximization in renewable dominated short-term electricity markets,” *International Transactions on Electrical Energy Systems (ITEES)*, vol. 3, 12 2021.
- [43] H. Haes Alhelou, M. E. Hamedani-Golshan, T. C. Njenda, and P. Siano, “A survey on power system blackout and cascading events: Research motivations and challenges,” *Energies*, vol. 12, no. 4, 2019.
- [44] ENTSO-E, *ENTSO-E Transmission System Map*. 2021. [Online]. Available: <https://www.entsoe.eu/data/map/> (visited on 10/12/2021).
- [45] L. Meeus, *The evolution of electricity markets in Europe*. Cheltenham, UK: Edward Elgar Publishing, 2020.

- [46] T. Jamasb and M. Pollitt, “Electricity market reform in the European Union: Review of progress towards liberalisation and integration,” *Energy J.*, vol. 26, pp. 11–41, 2005.
- [47] European Commission, *DG competition report on energy sector inquiry*. Luxembourg: Publications Office of the European Union, 2007.
- [48] TSCNET Services, *ENTSO-E: Conference on the "System of Systems"*. 2020. [Online]. Available: <https://www.tscnet.eu/entso-e-conference-on-the-system-of-systems/> (visited on 10/12/2021).
- [49] European Union Agency for the Cooperation of Energy Regulators (ACER) / Council of European Energy Regulators (CEER), *Annual Report on the Results of Monitoring the Internal Electricity and Natural Gas Markets in 2019, Electricity Wholesale Markets Volume*. 2020.
- [50] European Union Emission Trading Scheme, *Forward electricity market*. 2021. [Online]. Available: <https://www.emissions-euets.com/internal-electricity-market-glossary/1472-forward-electricity-market> (visited on 10/12/2021).
- [51] European Commission, *A guideline on capacity allocation and congestion management - 2015/1222 of 24 July 2015*. 2015.
- [52] R. A. van der Veen, A. Abbasy, and R. A. Hakvoort, “Agent-based analysis of the impact of the imbalance pricing mechanism on market behavior in electricity balancing markets,” *Energy Econ.*, vol. 34, no. 4, pp. 874–881, 2012.
- [53] ENTSO-E, *Electricity Balancing in Europe, An Overview of the European Balancing Market and electricity balancing guideline*. 2018.
- [54] ———, *Frequency Stability Evaluation Criteria for the Synchronous Zone of Continental Europe*. 2016. [Online]. Available: https://eepublicdownloads.entsoe.eu/clean-documents/SOC%20documents/RGCE_SPD_frequency_stability_criteria_v10.pdf (visited on 11/12/2021).
- [55] P. Cramton, “Electricity market design,” *Oxf. Rev. Econ. Policy*, vol. 33, no. 4, 2017.
- [56] L. Exizidis, *Electricity Markets with High Wind Power Penetration: Information Sharing and Incentive-Compatibility*. PhD thesis UMons.
- [57] D. Schönheit, M. Kenis, L. Lorenz, D. Möst, E. Delarue, and K. Bruninx, “Toward a fundamental understanding of flow-based market coupling for cross-border electricity trading,” *Adv. Appl. Energy*, vol. 2, 2021.
- [58] T. Kristiansen, “The flow based market coupling arrangement in europe: Implications for traders,” *Energy Strategy Rev.*, vol. 27, 2020.

Bibliography

- [59] M. Madani, *Revisiting European day-ahead electricity market auctions: MIP models and algorithms*. PhD thesis UCL, 2017.
- [60] Office of Gas and Electricity Markets (Ofgem), *Bidding Zones Literature Review*. [Online]. Available: https://www.ofgem.gov.uk/sites/default/files/docs/2014/10/fta_bidding_zone_configuration_literature_review_1.pdf (visited on 10/12/2021).
- [61] D. Newbery, G. Strbac, and I. Viehoff, “The benefits of integrating European electricity markets,” *Energy Policy*, vol. 94, pp. 253–263, 2016.
- [62] U. Gunturu and A. Schlosser, “Behavior of the aggregate wind resource in the ISO regions in the United States,” *Appl. Energy*, vol. 144, pp. 175–181, 2015.
- [63] European Commission, *CASE AT.39952 – POWER EXCHANGES*. Luxembourg: Publications Office of the European Union, 2014.
- [64] Nominated Electricity Market Operator (NEMO) committee, *EUPHEMIA Public Description, Single Price Coupling Algorithm*. 2019. [Online]. Available: https://www.nemo-committee.eu/assets/files/190410_Euphemia%20Public%20Description%20version%20NEMO%20Committee.pdf (visited on 10/12/2021).
- [65] J. Arroyo and A. Conejo, “Optimal response of a thermal unit to an electricity spot market,” *IEEE Transact. Power Syst.*, vol. 15, no. 3, pp. 1098–1104, 2000.
- [66] A. Conejo, F. Nogales, and J. Arroyo, “Price-taker bidding strategy under price uncertainty,” *IEEE Trans. Power Sys.*, vol. 17, no. 4, pp. 1081–1088, 2002.
- [67] T. Hong, P. Pinson, Y. Wang, R. Weron, D. Yang, and H. Zareipour, “Energy forecasting: A review and outlook,” *IEEE Open Access Journal of Power and Energy*, vol. 7, pp. 376–388, 2020.
- [68] P. Pinson, C. Chevallier, and G. N. Kariniotakis, “Trading wind generation from short-term probabilistic forecasts of wind power,” *IEEE Trans. Power Syst.*, vol. 22, no. 3, pp. 1148–1156, 2007.
- [69] A. Conejo, M. Carrion, and J. Morales, *Decision Making Under Uncertainty in Electricity Markets*, ser. Operations Research and Management Science. New York: Springer, 2010.
- [70] J.-F. Toubeau, *Stochastic optimization of Virtual Power Plants participating in electricity markets: from forecasting to decision-making*. PhD thesis UMons, 2018.

- [71] J. van der Hertten, F. Depuydt, L. De Baets, D. Deschrijver, M. Strobbe, C. Develder, T. Dhaene, R. Bruneliere, and J.-W. Rombouts, “Energy flexibility assessment of an industrial coldstore process,” in *2016 IEEE International Energy Conference (ENERGYCON)*, 2016, pp. 1–6. DOI: 10.1109/ENERGYCON.2016.7514032.
- [72] K. Bruninx, H. Pandžić, H. Le Cadre, and E. Delarue, “On the interaction between aggregators, electricity markets and residential demand response providers,” *IEEE Transact. Power Syst.*, vol. 35, no. 2, pp. 840–853, 2020.
- [73] Z. Liu, Q. Wu, S. Huang, L. Wang, M. Shahidehpour, and Y. Xue, “Optimal day-ahead charging scheduling of electric vehicles through an aggregative game model,” *IEEE Transact. Smart Grid*, vol. 9, no. 5, pp. 5173–5184, 2018.
- [74] M. Zugno, J. Morales, P. Pinson, and H. Madsen, “Pool strategy of a price-maker wind power producer,” *IEEE Trans. Power Syst.*, vol. 28, no. 3, pp. 3440–3450, 2013.
- [75] M. Zima-Bočkarjova, J. Matevosyan, M. Zima, and L. Söder, “Sharing of profit from coordinated operation planning and bidding of hydro and wind power,” *IEEE Transact. Power Syst.*, vol. 25, no. 3, pp. 1663–1673, 2010.
- [76] J. M. Morales, A. J. Conejo, and J. PÉrez-Ruiz, “Short-term trading for a wind power producer,” *IEEE Transact. Power Syst.*, vol. 25, no. 1, pp. 554–564, 2010.
- [77] L. H. Lam, V. Ilea, and C. Bovo, “European day-ahead electricity market coupling: Discussion, modeling, and case study,” *Electric Power Systems Research*, vol. 155, pp. 80–92, 2018.
- [78] Y. Ye, D. Papadaskalopoulos, J. Kazempour, and G. Strbac, “Incorporating non-convex operating characteristics into bi-level optimization electricity market models,” *IEEE Transact. Power Syst.*, vol. 35, no. 1, pp. 163–176, 2020.
- [79] M. Karasavvidis, *Offering Strategy of a Power Producer Using Block Offers*. M.Sc. Thesis Technical University of Denmark, 2019.
- [80] G. De Viviero-Serrano, K. Bruninx, and E. Delarue, “Implications of bid structures on the offering strategies of merchant energy storage systems,” *Appl. Energy*, vol. 251, p. 113375, 2019, ISSN: 0306-2619.
- [81] Y. Zhou, C. Wang, J. Wu, J. Wang, M. Cheng, and G. Li, “Optimal scheduling of aggregated thermostatically controlled loads with renewable generation in the intraday electricity market,” *Applied Energy*, vol. 188, pp. 456–465, 2017.

Bibliography

- [82] D. Wozabal and G. Rameseder, “Optimal bidding of a virtual power plant on the spanish day-ahead and intraday market for electricity,” *European Journal of Operational Research*, vol. 280, no. 2, pp. 639–655, 2020.
- [83] H. T. Nguyen, L. B. Le, and Z. Wang, “A bidding strategy for virtual power plants with the intraday demand response exchange market using the stochastic programming,” *IEEE Transact. Ind. Appl.*, vol. 54, no. 4, pp. 3044–3055, 2018.
- [84] S. Yao, P. Wang, X. Liu, H. Zhang, and T. Zhao, “Rolling optimization of mobile energy storage fleets for resilient service restoration,” *IEEE Transact. Smart Grid*, vol. 11, no. 2, pp. 1030–1043, 2020.
- [85] G. Bertrand and A. Papavasiliou, “Adaptive trading in continuous intraday electricity markets for a storage unit,” *IEEE Transact. Power Syst.*, vol. 35, no. 3, 2020.
- [86] I. Boukas, D. Ernst, T. Théate, A. Bolland, A. Huynen, M. Buchwald, C. Wynants, and B. Cornélusse, “A deep reinforcement learning framework for continuous intraday market bidding,” *Machine Learning*, 2021.
- [87] P. Shinde, I. Kouveliotis-Lysikatos, and M. Amelin, “Multistage stochastic programming for vpp trading in continuous intraday electricity markets,” *IEEE Transact. Sust. Energy*, pp. 1–1, 2022.
- [88] P. Zou, Q. Chen, Q. Xia, G. He, and C. Kang, “Evaluating the contribution of energy storages to support large-scale renewable generation in joint energy and ancillary service markets,” *IEEE Transact. Sust. Energy*, vol. 7, no. 2, pp. 808–818, 2016.
- [89] J. E. Contreras-Ocana, M. A. Ortega-Vazquez, and B. Zhang, “Participation of an energy storage aggregator in electricity markets,” *IEEE Transact. Smart Grid*, vol. 10, no. 2, pp. 1171–1183, 2019.
- [90] J.-F. Toubreau, J. Bottieau, Z. De Greve, F. Vallee, and K. Bruninx, “Data-driven scheduling of energy storage in day-ahead energy and reserve markets with probabilistic guarantees on real-time delivery,” *IEEE Trans. Power Syst.*, 2020.
- [91] G. He, Q. Chen, C. Kang, P. Pinson, and Q. Xia, “Optimal bidding strategy of battery storage in power markets considering performance-based regulation and battery cycle life,” *IEEE Transact. Smart Grid*, vol. 7, no. 5, pp. 2359–2367, 2016.
- [92] E. Nasrolahpour, J. Kazempour, H. Zareipour, and W. D. Rosehart, “A bilevel model for participation of a storage system in energy and reserve markets,” *IEEE Transact. Sust. Energy*, vol. 9, no. 2, pp. 582–598, 2018.

- [93] K. Pandžić, K. Bruninx, and H. Pandžić, “Managing risks faced by strategic battery storage in joint energy-reserve markets,” *IEEE Transact. Power Systems*, vol. 36, no. 5, pp. 4355–4365, 2021.
- [94] A. Papavasiliou, “Scarcity pricing and the missing european market for real-time reserve capacity,” *The Electricity Journal*, vol. 33, no. 10, 2020.
- [95] Elia Group, *Data Download*. 2019. [Online]. Available: [%5Curl%7Bwww.elia.be/en/grid-data/data-download%7D,%20urldate%20=%20%7B2021-10-12%7D](https://www.elia.be/en/grid-data/data-download).
- [96] J. Lago, F. De Ridder, and B. De Schutter, “Forecasting spot electricity prices: Deep learning approaches and empirical comparison of traditional algorithms,” *Applied Energy*, vol. 221, pp. 386–405, 2018.
- [97] B. Lim and S. Zohren, “Time-series forecasting with deep learning: A survey,” *Phil. Trans. R. Soc. A.*, vol. 379, no. 2194, p. 20 200 209, 2021.
- [98] T. Hong, P. Pinson, and S. Fan, “Global energy forecasting competition 2012,” *Int. J. Forecast.*, vol. 30, no. 2, pp. 357–363, 2014.
- [99] Y. Zhang, J. Wang, and X. Wang, “Review on probabilistic forecasting of wind power generation,” *Renewable & Sustainable Energy Reviews*, vol. 32, no. 5, pp. 255–270, 2014.
- [100] J. Nowotarski and R. Weron, “Recent advances in electricity price forecasting: A review of probabilistic forecasting,” *Renewable & Sustainable Energy Reviews*, vol. 81, pp. 1548–1568, 2018.
- [101] T. Hong and S. Fan, “Probabilistic electric load forecasting: A tutorial review,” *Int. J. Forecast.*, 2016.
- [102] J. Juban, L. Fugon, and G. Kariniotakis, “Probabilistic short-term wind power forecasting based on kernel density estimators,” *Proc. Eur. Wind Energy Conf.*, pp. 1–11, 2007.
- [103] A. Khosravi, S. Nahavandi, and D. Creighton, “Construction of optimal prediction intervals for load forecasting problem,” *IEEE Trans. Power Syst.*, vol. 25, no. 3, pp. 1496–1503, 2010.
- [104] C. Wan, Z. Xu, P. Pinson, Z. Y. Dong, and K. P. Wong, “Probabilistic forecasting of wind power generation using extreme learning machine,” *IEEE Trans. Power Syst.*, vol. 29, no. 3, pp. 1033–1044, 2014.
- [105] F. Golestaneh, P. Pinson, and H. B. Gooi, “Very short-term nonparametric probabilistic forecasting of renewable energy generation - with application to solar energy,” *IEEE Trans. Power Syst.*, vol. 31, no. 5, pp. 3850–386. 2016.

Bibliography

- [106] C. Wan, J. Lin, Y. Song, Z. Xu, and G. Yang, “Probabilistic forecasting of photovoltaic generation: An efficient statistical approach,” *IEEE Trans. Power Syst.*, vol. 32, no. 3, pp. 2471–2472, 2017.
- [107] Z. Shi, H. Liang, and V. Dinavahi, “Direct interval forecast of uncertain wind power based on recurrent neural networks,” *IEEE Trans. Sust. Energy*, vol. 9, no. 3, 2018.
- [108] J. Dowell and P. Pinson, “Very-short-term probabilistic wind power forecasts by sparse vector autoregression,” *IEEE Trans. Smart Grid*, vol. 7, no. 2, pp. 763–770, 2016.
- [109] A. Kavousi-Fard, A. Khosravi, and S. Nahavandi, “A new fuzzy-based combined prediction interval for wind power forecasting,” *IEEE Trans. Power Syst.*, vol. 31, no. 1, pp. 18–26, 2016.
- [110] R. Bessa, V. Miranda, A. Botterud, J. Wang, and E. Constantinescu, “Time adaptive conditional kernel density estimation for wind power forecasting,” *IEEE Trans. Sust. Energy*, vol. 3, no. 4, pp. 660–669, 2012.
- [111] Y. Zhang and J. Wang, “K-nearest neighbors and a kernel density estimator for gefcom2014 probabilistic wind power forecasting,” *Int. J. Forecasting*, vol. 32, no. 3, pp. 1074–1080, 2016.
- [112] C. Wan, J. Lin, J. Wang, Y. Song, and Z. Y. Dong, “Direct quantile regression for nonparametric probabilistic forecasting of wind power generation,” *IEEE Trans. Power Syst.*, vol. 32, no. 4, pp. 2767–2778, 2017.
- [113] K. Hatalis, A. Lamadrid, K. Scheinberg, and S. Kishore, “Smooth pin-ball neural network for probabilistic forecasting of wind power,” *arXiv preprint arXiv:1710.01720*, 2017.
- [114] M. Khodayar and J. Wang, “Spatio-temporal graph deep neural network for short-term wind speed forecasting,” *IEEE Trans. Sust. Energy*, vol. 10, no. 2, pp. 670–681, 2019.
- [115] Elia Group, *Tariffs for maintaining and restoring the residual balance of the individual access responsible parties 2016-2019*. 2021. [Online]. Available: <https://www.elia.be/en/customers/invoicing-and-tariffs> (visited on 10/12/2021).
- [116] A. F. Zambrano and L. F. Giraldo, “Solar irradiance forecasting models without on-site training measurements,” *Renew. Energy*, vol. 152, pp. 557–566, 2020.
- [117] C. Nwankpa, W. L. Ijomah, A. Gachagan, and S. Marshall, “Activation functions: Comparison of trends in practice and research for deep learning,” *arXiv preprint arXiv:1811.03378*, 2018.

- [118] J. Toubeau, J. Bottieau, F. Vallée, and Z. De Grève, “Deep learning-based multivariate probabilistic forecasting for short-term scheduling in power markets,” *IEEE Trans Power Syst.*, vol. 34, no. 2, pp. 1203–1215, 2019.
- [119] Y. Wang, N. Zhang, Q. Chen, D. S. Kirschen, P. Li, and Q. Xia, “Data-driven probabilistic net load forecasting with high penetration of behind-the-meter pv,” *IEEE Transact. Power Syst.*, vol. 33, no. 3, pp. 3255–3264, 2018.
- [120] A. Graves, M. Liwicki, S. Fernández, R. Bertolami, H. Bunke, and J. Schmidhuber, “A novel connectionist system for unconstrained handwriting recognition,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 31, no. 5, pp. 855–868, 2009.
- [121] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [122] A. Graves and J. Schmidhuber, “Framewise phoneme classification with bidirectional lstm and other neural network architectures,” *Neural Netw.*, vol. 18, no. 5-6, pp. 602–10, 2005.
- [123] K. Cho, B. van Merriënboer, D. Bahdanau, and Y. Bengio, “On the properties of neural machine translation: Encoder–decoder approaches,” in *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, 2014, pp. 103–111.
- [124] K. C. D. Bahdanau and Y. Bengio, “Neural machine translation by jointly learning to align and translate,” *arXiv preprint arXiv:1409.0473*, 2014.
- [125] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2017.
- [126] M. Abadi and A. Agarwal et al., “Tensorflow: Large-scale machine learning on heterogeneous distributed systems,” *arXiv preprint*, 1603.04467, 2016.
- [127] J. Bergstra and Y. Bengio, “Random search for hyper-parameter optimization,” *Journal of machine learning research*, vol. 13, no. 2, 2012.
- [128] Y. Wang, D. Gan, M. Sun, N. Zhang, Z. Lu, and C. Kang, “Probabilistic individual load forecasting using pinball loss guided lstm,” *Applied Energy*, pp. 10–20, 2019.
- [129] Y. Wang, N. Zhang, Y. Tan, T. Hong, D. S. Kirschen, and C. Kang, “Combining probabilistic load forecasts,” *IEEE Trans. Smart Grid*, vol. 10, no. 4, pp. 3664–3674, 2019.

Bibliography

- [130] P. Pinson, “Very-short-term probabilistic forecasting of wind power with generalized logit—normal distributions,” *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, vol. 61, pp. 555–576, 2012.
- [131] S. Seabold and J. Perktold, “Statsmodels: Econometric and statistical modeling with python,” in *9th Python in Science Conference*, 2010.
- [132] N. Meinshausen, “Quantile regression forests,” *Journal of Machine Learning Research*, pp. 983–999, 2006.
- [133] J. H. Friedman, “Greedy function approximation: A gradient boosting machine.,” *The Annals of Statistics*, vol. 29, no. 5, pp. 1189–1232, 2021.
- [134] T. Gneiting, F. Balabdaoui, and A. E. Raftery, “Probabilistic forecasts, calibration and sharpness,” *J. Royal Stat. Soc. B*, pp. 243–268, 2007.
- [135] R. Philipsen, G. Morales-España, M. de Weerd, and L. de Vries, “Trading power instead of energy in day-ahead electricity markets,” *Appl. Energy*, vol. 233-234, pp. 802–815, 2019.
- [136] J. H. Verberk, R. M. Hermans, P. P. J. Van den Bosch, A. Jokić, and J. Frunt, “Systematic design of market-based balancing arrangements for deregulated power systems: An asynchronous solution,” in *2011 IEEE Trondheim PowerTech*, 2011, pp. 1–7.
- [137] TenneT, *Imbalance Management TenneT, Analysis report*. 2011.
- [138] C. Koch and P. Maskos, “Passive balancing through intraday trading: Whether interactions between short-term trading and balancing stabilize germany’s electricity system,” *IJEEP*, vol. 10, pp. 101–112, 2020.
- [139] Commission of Regulation of Electricity and Gas (CREG), *Study on the functioning and price evolution of the Belgian wholesale electricity market, Monitoring Report 2020*. 2021.
- [140] —, *Study on the functioning and price evolution of the Belgian wholesale electricity market, Monitoring Report 2017*. 2018.
- [141] European Federation of Energy Traders (EFET), *EFET position paper, Price formation and capacity withholding in light of Regulation (EU) 2019/943 and Regulation (EU) 1227/2011*. 2020.
- [142] Elia Group, *Study report on Scarcity Pricing in the context of the 2018 discretionary incentives*. 2021. [Online]. Available: <https://www.elia.be/fr/marche-de-electricite-et-reseau/studies/scarcity-pricing-simulation> (visited on 10/12/2021).
- [143] J. Bottieau, L. Hubert, Z. De Grève, F. Vallée, and J.-F. Toubéau, “Very-short-term probabilistic forecasting for a risk-aware participation in the single price imbalance settlement,” *IEEE Trans. Power Syst.*, vol. 35, no. 2, pp. 1218–1230, 2020.

- [144] S. A. Gabriel, A. J. Conejo, J. D. Fuller, B. F. Hobbs, and C. Ruiz, *Complementarity modeling in energy markets*, ser. Springer Science and Business Media. New York: Springer, 2012.
- [145] J. Fortuny-Amat and B. McCarl, “A representation and economic interpretation of a two-level programming problem,” *Journal of the Operational Research Society*, vol. 32, no. 9, pp. 783–792, 1981.
- [146] S. Pineda and J. Morales, “Solving linear bilevel problems using big-ms: Not all that glitters is gold,” *IEEE Trans. Power Syst.*, vol. 34, no. 3, pp. 2469–2471, 2019.
- [147] H. Quan, D. Srinivasan, and A. Khosravi, “Incorporating wind power forecast uncertainties into stochastic unit commitment using neural network-based prediction intervals,” *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 26, no. 9, pp. 2123–2135, 2015.
- [148] M. Kazemi, H. Zareipour, M. Ehsan, and W. D. Rosehart, “A robust linear approach for offering strategy of a hybrid electric energy company,” *IEEE Trans. Power Syst.*, vol. 32, no. 3, pp. 1949–1959, 2017.
- [149] J. M. Morales, A. J. Conejo, H. Madsen, P. Pinson, and M. Zugno, *Integrating Renewables in Electricity Markets: Operational Problems*, ser. Operations Research and Management Science. New York: Springer, 2014.
- [150] R. Dahlgren, Chen-Ching Liu, and J. Lawarree, “Risk assessment in energy trading,” *IEEE Trans. Power Syst.*, vol. 18, no. 2, pp. 503–511, 2003.
- [151] K. Zare, A. J. Conejo, M. Carrión, and M. P. Moghaddam, “Multi-market energy procurement for a large consumer using a risk-aversion procedure,” *Electr. Pow. Syst. Res.*, vol. 80, no. 1, pp. 63–70, 2010.
- [152] M. Kazemi, B. Mohammadi-Ivatloo, and M. Ehsan, “Risk-constrained strategic bidding of gencos considering demand response,” *IEEE Trans. Power Syst.*, vol. 30, no. 1, pp. 376–384, 2015.
- [153] M. Shabanzadeh, M.-K. Sheikh-El-Eslami, and M.-R. Haghifam, “The design of a risk-hedging tool for virtual power plants via robust optimization approach,” *Applied Energy*, vol. 155, pp. 766–777, 2015, ISSN: 0306-2619.
- [154] T. Dai and W. Qiao, “Optimal bidding strategy of a strategic wind power producer in the short-term market,” *IEEE Trans. Sust. Energy*, vol. 6, no. 3, pp. 707–719, 2015.
- [155] L. Baringo and A. J. Conejo, “Offering strategy of wind-power producer: A multi-stage risk-constrained approach,” *IEEE Trans. Power Syst.*, vol. 31, no. 2, pp. 1420–1429, 2016, ISSN: 1558-0679.

Bibliography

- [156] E. G. Kardakos, C. K. Simoglou, and A. G. Bakirtzis, "Optimal offering strategy of a virtual power plant: A stochastic bi-level approach," *IEEE Trans. Smart Grid*, vol. 7, no. 2, pp. 794–806, 2016.
- [157] T. Li, M. Shahidehpour, and Z. Li, "Risk-constrained bidding strategy with stochastic unit commitment," *IEEE Trans. Power Syst.*, vol. 22, no. 1, pp. 449–458, 2007.
- [158] P. Sheikhhahmadi, S. Bahramara, J. Moshtagh, and M. Yazdani Damavandi, "A risk-based approach for modeling the strategic behavior of a distribution company in wholesale energy market," *Appl. Energy*, vol. 214, pp. 24–38, 2018, ISSN: 0306-2619.
- [159] N. Shang, C. Ye, Y. Ding, T. Tu, and B. Huo, "Risk-based optimal power portfolio methodology for generation companies considering cross-region generation right trade," *Appl. Energy*, vol. 254, p. 113511, 2019, ISSN: 0306-2619.
- [160] Y. Wang, Y. Dvorkin, R. Fernández-Blanco, B. Xu, T. Qiu, and D. S. Kirschen, "Look-ahead bidding strategy for energy storage," *IEEE Trans. Sustain. Energy*, vol. 8, no. 3, pp. 1106–1117, 2017.
- [161] U. Yıldıran and İ. Kayahan, "Risk-averse stochastic model predictive control-based real-time operation method for a wind energy generation system supported by a pumped hydro storage unit," *Appl. Energy*, vol. 226, pp. 631–643, 2018, ISSN: 0306-2619.
- [162] V. Guerrero-Mestre, A. A. Sánchez de la Nieta, J. Contreras, and J. P. S. Catalão, "Optimal bidding of a group of wind farms in day-ahead markets through an external agent," *IEEE Trans. Power Syst.*, vol. 31, no. 4, pp. 2688–2700, Sep. 2016.
- [163] A. A. Sánchez de la Nieta, N. G. Paterakis, and M. Gibescu, "Participation of photovoltaic power producers in short-term electricity markets based on rescheduling and risk-hedging mapping," *Appl. Energy*, vol. 266, p. 114741, 2020, ISSN: 0306-2619.
- [164] M. Song and M. Amelin, "Price-maker bidding in day-ahead electricity market for a retailer with flexible demands," *IEEE Trans. Power Syst.*, vol. 33, no. 2, pp. 1948–1958, 2018.
- [165] H. Rashidizadeh-Kermani, M. Vahedipour-Dahraie, M. Shafie-khah, and P. Siano, "A regret-based stochastic bi-level framework for scheduling of DR aggregator under uncertainties," *IEEE Trans. Smart Grid*, vol. 11, no. 4, pp. 3171–3184, 2020.
- [166] K. Bruninx, H. Pandžić, H. Le Cadre, and E. Delarue, "On the interaction between aggregators, electricity markets and residential demand response providers," *IEEE Trans. Power Syst.*, vol. 35, no. 2, pp. 840–853, 2020.

- [167] M. S. Al-Swaiti, A. T. Al-Awami, and M. W. Khalid, “Co-optimized trading of wind-thermal-pumped storage system in energy and regulation markets,” *Energy*, vol. 138, pp. 991–1005, 2017, ISSN: 0360-5442.
- [168] M. H. Abbasi, M. Taki, A. Rajabi, L. Li, and J. Zhang, “Coordinated operation of electric vehicle charging and wind power generation as a virtual power plant: A multi-stage risk constrained approach,” *Appl. Energy*, vol. 239, pp. 1294–1307, 2019, ISSN: 0306-2619.
- [169] Y. Liu, Z. Shen, X. Tang, H. Lian, J. Li, and J. Gong, “Worst-case conditional value-at-risk based bidding strategy for wind-hydro hybrid systems under probability distribution uncertainties,” *Appl. Energy*, vol. 256, p. 113918, 2019, ISSN: 0306-2619.
- [170] S. Ghavidel, M. J. Ghadi, A. Azizivahed, J. Aghaei, L. Li, and J. Zhang, “Risk-constrained bidding strategy for a joint operation of wind power and caes aggregators,” *IEEE Trans. Sust. Energy*, vol. 11, no. 1, pp. 457–466, 2020.
- [171] H. Höschle, H. Le Cadre, Y. Smeers, A. Papavasiliou, and R. Belmans, “An admm-based method for computing risk-averse equilibrium in capacity markets,” *IEEE Trans. Power Syst.*, vol. 33, no. 5, pp. 4819–4830, 2018.
- [172] J.-F. Toubeau, Z. De Grève, and F. Vallée, “Medium-term multimarket optimization for virtual power plants: A stochastic-based decision environment,” *IEEE Trans. Power Syst.*, vol. 33, no. 2, pp. 1399–1410, 2018.
- [173] T. Hong, J. Xie, and J. Black, “Global energy forecasting competition 2017: Hierarchical probabilistic load forecasting,” *Int. J. Forecast.*, vol. 35, pp. 1389–1399, 2019. DOI: 10.1016/j.ijforecast.2019.02.006.
- [174] G. Dalal, E. Gilboa, S. Mannor, and L. Wehenkel, “Unit commitment using nearest neighbor as a short-term proxy,” *2018 PSCC*, pp. 1–7, 2018.
- [175] L. Duchesne, E. Karangelos, and L. Wehenkel, “Machine learning of real-time power systems reliability management response,” *2017 IEEE Manchester PowerTech*, pp. 1–6, 2017.
- [176] G. Dalal, E. Gilboa, S. Mannor, and L. Wehenkel, “Chance-constrained outage scheduling using a machine learning proxy,” *IEEE Trans. Power Syst.*, vol. 34, no. 4, pp. 2528–2540, 2019.
- [177] J.-F. Toubeau, P.-D. Dapoz, J. Bottieau, A. Wautier, Z. De Grève, and F. Vallée, “Recalibration of recurrent neural networks for short-term wind power forecasting,” *Electr. Power Syst. Res.*, vol. 190, p. 106639, 2021, ISSN: 0378-7796.

Bibliography

- [178] M. Du, N. Liu, and X. Hu, “Techniques for interpretable machine learning,” *Commun. ACM*, vol. 63, no. 1, pp. 68–77, 2019, ISSN: 0001-0782.
- [179] R. Cynthia, “Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead.,” *Nat. Mach. Intell.*, vol. 1, pp. 206–215, 2019.
- [180] A. A. Ismail, M. Gunady, H. C. Bravo, and S. Feizi, “Benchmarking deep learning interpretability in time series predictions,” in *Adv. Neural Inf. Process. Syst.*, vol. 33, 2020.
- [181] Z. C. Lipton, “The mythos of model interpretability,” *Queue*, vol. 16, no. 3, pp. 31–57, 2018.
- [182] F. Doshi-Velez and B. Kim, “Towards a rigorous science of interpretable machine learning,” arXiv preprint, 1702.08608, 2017.
- [183] Y. Zhang, P. Tiño, A. Leonardis, and K. Tang, “A survey on neural network interpretability,” *IEEE Trans. Emerg. Top. Comput. Intell.*, vol. 5, no. 5, pp. 726–742, 2021.
- [184] W. Samek, G. Montavon, S. Lapuschkin, C. J. Anders, and K.-R. Müller, “Explaining deep neural networks and beyond: A review of methods and applications,” *Proceedings of the IEEE*, vol. 109, no. 3, pp. 247–278, 2021.
- [185] Z. Yang, A. Zhang, and A. Sudjianto, “Enhancing explainability of neural networks through architecture constraints,” *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 32, no. 6, pp. 2610–2621, 2021.
- [186] B. Lim, S. O. Arik, N. Loeff, and T. Pfister, “Temporal fusion transformers for interpretable multi-horizon time series forecasting,” arXiv preprint, 1912.09363, 2020.
- [187] M. T. Ribeiro, S. Singh, and C. Guestrin, “Why should i trust you? explaining the predictions of any classifier,” *22nd ACM SIGKDD international conference on knowledge discovery and data mining*, 2016.
- [188] S. Lundberg and S.-I. Lee, “A unified approach to interpreting model predictions,” *Advances in Neural Information Processing Systems*, pp. 4765–4774, 2017.
- [189] Y. Lu, I. Murzakhanov, and S. Chatzivasileiadis, “Neural network interpretability for forecasting of aggregated renewable generation,” in *2021 IEEE International Conference on SmartGridComm*, 2021, pp. 282–288.
- [190] T. Gale, E. Elsen, and S. Hooker, “The state of sparsity in deep neural networks,” arXiv preprint, 1902.09574, 2019.

- [191] D. W. Bunn, J. N. Inekwe, and D. MacGeehan, “Analysis of the fundamental predictability of prices in the british balancing market,” *IEEE Trans. Power Syst.*, vol. 36, no. 2, pp. 1309–1316, 2021.
- [192] R. Weron, “Electricity price forecasting: A review of the state-of-the-art with a look into the future,” *Int J Forecast*, vol. 30, pp. 1030–1081, 4 2014, ISSN: 0169-2070.
- [193] J. Lago, G. Marcjasz, B. De Schutter, and R. Weron, “Forecasting day-ahead electricity prices: A review of state-of-the-art algorithms, best practices and an open-access benchmark,” *Appl. Energy*, vol. 293, p. 116983, 2021, ISSN: 0306-2619.
- [194] M. Olsson and L. Soder, “Modeling real-time balancing power market prices using combined sarima and markov processes,” *IEEE Trans. Power Syst.*, vol. 23, no. 2, pp. 443–450, 2008.
- [195] I. Dimoukias, M. Amelin, and M. R. Hesamzadeh, “Forecasting balancing market prices using hidden markov models,” in *2016 13th International Conference on the EEM*, 2016, pp. 1–5.
- [196] M. B. Olsson and L. Söder, “Modeling swedish real-time balancing power prices using nonlinear time series models,” in *2010 IEEE 11th International Conference on PMAPS*, 2010, pp. 358–363.
- [197] G. Klaeboe, A. L. Eriksrud, and S.-E. Fleten, “Benchmarking time series based forecasting models for electricity balancing market prices,” *Energy Syst.*, vol. 6, pp. 43–61, 2015.
- [198] S. Jaehnert, H. Farahmand, and G. L. Doorman, “Modelling of prices using the volume in the norwegian regulating power market,” in *2009 IEEE Bucharest PowerTech*, 2009, pp. 1–7.
- [199] T. Jónsson, P. Pinson, H. A. Nielsen, and H. Madsen, “Exponential smoothing approaches for prediction in real-time electricity markets,” *Energies*, vol. 7, no. 6, pp. 3710–3732, 2014, ISSN: 1996-1073.
- [200] A. Lucas, K. Pegios, E. Kotsakis, and D. Clarke, “Price Forecasting for the Balancing Energy Market Using Machine-Learning Regression,” *Energies*, vol. 13, no. 20, pp. 1–16, 2020.
- [201] D. He and W.-P. Chen, “A real-time electricity price forecasting based on the spike clustering analysis,” in *2016 IEEE/PES T & D*, 2016, pp. 1–5.
- [202] H. Yang and K. R. Schell, “Hfnet: Forecasting real-time electricity price via novel gru architectures,” in *2020 International Conference on PMAPS*, 2020, pp. 1–6.

Bibliography

- [203] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Adv. Neural Inf. Process. Syst.*, 2017, pp. 6000–6010.
- [204] A. Gillioz, J. Casas, E. Mugellini, and O. A. Khaled, “Overview of the transformer-based models for nlp tasks,” in *2020 15th Conference on FedCSIS*, 2020, pp. 179–183.
- [205] T. Q. Nguyen and J. Salazar, “Transformers without tears: Improving the normalization of self-attention,” *IWSLT*, 2019.
- [206] B. Lim and S. Zohren, “Time-series forecasting with deep learning: A survey,” *Phil. Trans. R. Soc. A.*, vol. 379, no. 2194, p. 20 200 209, 2021, ISSN: 1471-2962.
- [207] D.-A. Clevert, T. Unterthiner, and S. Hochreiter, “Fast and accurate deep network learning by exponential linear units (elus),” arXiv preprint, 1511.07289, 2016.
- [208] C. Guo and F. Berkhahn, “Entity embeddings of categorical variables,” arXiv preprint, 1604.06737, 2016.
- [209] S. Li, X. Jin, Y. Xuan, X. Zhou, W. Chen, Y.-X. Wang, and X. Yan, “Enhancing the locality and breaking the memory bottleneck of transformer on time series forecasting,” arXiv preprint, 1907.00235, 2020.
- [210] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” arXiv preprint, 1512.03385, 2015.

APPENDIX A.

List of Publications

In reviewing process

Peer-review journal articles

- **J. Bottieau**, Y. Wang, Z. De Grève, F. Vallée, and J-F. Toubeau, "Transformer Model for Interpretable Probabilistic Forecasting of Real-Time Electricity Prices," currently in second round of review in IEEE Trans. Power Syst.
- **J. Bottieau**, Z. De Grève, T. Piraux, A. Dubois, F. Vallée, and J-F. Toubeau, "A Cross-Learning Approach for Cold-Start Forecasting of Residential Photovoltaic Generation," forthcoming in Electric Power Systems Research.

Peer-review Conference papers

- **J. Bottieau**, Z. De Grève, T. Piraux, A. Dubois, F. Vallée, and J-F. Toubeau, "A Cross-Learning Approach for Cold-Start Forecasting of Residential Photovoltaic Generation," forthcoming in 2022 IEEE Power Systems Computation Conference (PSCC), Porto, Portugal

Publications in Electrical Engineering

Peer-review journal articles

- **J. Bottieau**, K. Bruninx, A. Sanjab, Z. De Grève, F. Vallée and J-F. Toubeau, "Automatic Risk Adjustment for Profit Maximization in

Appendix A. List of Publications

- Renewable Dominated Short-Term Electricity Markets," in ITEES, vol. 3, issue 12, 2021.
- J-F. Toubreau, **J. Bottieau**, Z. De Grève, F. Vallée, K. Bruninx, "Data-Driven Scheduling of Energy Storage in Day-Ahead Energy and Reserve Markets with Probabilistic Guarantees on Real-Time Delivery," in IEEE Transact. Power Syst., vol. 36, no. 4, pp. 2815-2828, 2021.
 - J-F. Toubreau, **J. Bottieau**, Y. Wang, and F. Vallée, "Interpretable Probabilistic Forecasting of Imbalances in Renewable-Dominated Electricity Systems," in IEEE Trans. Sust. Energy. Early access: <https://ieeexplore.ieee.org/abstract/document/9464660>.
 - J-F. Toubreau, T. Morstyn, **J. Bottieau**, K. Zheng, D. Apostolopoulou, Z. De Grève, Y. Wang, F. Vallée, "Capturing Spatio-Temporal Dependencies in the Probabilistic Forecasting of Distribution Locational Marginal Prices," in IEEE Transact. Smart Grid, vol. 12, no. 3, pp. 2663-2674, 2021.
 - J-F. Toubreau, P-D. Dapoz, **J. Bottieau**, A. Wautier, Z. De Grève, F. Vallée, "Recalibration of Recurrent Neural Networks for Short-Term Wind Power Forecasting," in Electric Power Systems Research, vol. 190, pp. 1-7, 2021.
 - Z. De Grève, **J. Bottieau**, D. Vangulick, A. Wautier, P-D. Dapoz, A. Arrigo, J-F. Toubreau, F. Vallée, "Machine Learning Techniques for Improving Self-Consumption in Renewable Energy Communities," in Energies, vol. 13, issue 18, pp. 4892-4909, 2020.
 - **J. Bottieau**, L. Hubert, Z. De Grève, F. Vallée, J-F. Toubreau, "Very Short-Term Probabilistic Forecasting for a Risk-Aware Participation in the Single Price Imbalance Settlement," in IEEE Transact. Power Syst., vol. 35, no. 2, pp. 1218-1230, 2020.
 - J-F. Toubreau, S. Iassinovski, E. Jean, J-Y. Parfait, **J. Bottieau**, Z. De Grève, F. Vallée, "A Nonlinear Hybrid Approach for the Scheduling of Merchant Underground Pumped Hydro Energy Storage," in IET Generation, Transmission and Distribution, vol. 13, issue 21, pp. 4798-4808, 2019.
 - J-F. Toubreau, **J. Bottieau**, F. Vallée, Z. De Grève, "Deep Learning-based Multivariate Probabilistic Forecasting for Short-Term Scheduling in Power Markets," in IEEE Transact. Power Syst., vol 34, no. 2, pp. 1203-1215, 2019.

Peer-review Conference papers

- **J. Bottieau**, Arrigo Adriano, Z. De Grève, F. Vallée, J-F. Toubeau, "A Distributionally Robust Framework for Improving the Provision of Passive Balancing Services," in 2021 IEEE PowerTech International Conference, Madrid, Spain, 2021.
- J-F. Toubeau, P-D. Dapoz, **J. Bottieau**, A. Wautier, Z. De Grève, F. Vallée, "Recalibration of Recurrent Neural Networks for Short-Term Wind Power Forecasting," in 2020 IEEE Power Systems Computation Conference (PSCC), Porto, Portugal, 2020.
- **J. Bottieau**, F. Vallée, Z. De Grève, J-F. Toubeau, "Leveraging Provision of Frequency Regulation Services from Wind Generation by Improving Day-Ahead Predictions using LSTM Neural Networks," in 2018 IEEE International Energy Conference (ENERGYCON), Limassol, Cyprus, 2018.
- **J. Bottieau**, Z. De Grève, F. Vallée, J-F. Toubeau, "Fostering the Participation of Battery Storage Systems in the Real-Time Balancing of Electrical Grids," in 2018 IEEE Young Researchers Symposium in Electrical Power Engineering (YRS), Brussels, Belgium, 2018.
- J-F. Toubeau, **J. Bottieau**, F. Vallée, Z. De Grève, "Improved Day-Ahead Predictions of Load and Renewable Generation by Optimally Exploiting Multi-Scale Dependencies," in 2017 IEEE Conference on Innovative Smart Grids Technologies (ISGT Asia), Auckland, New Zealand (2017).

Publications outside Electrical Engineering

Peer-review journal articles

- C. Liefferinckx, **J. Bottieau**, J-F. Toubeau, D. Thomas, J-F. Rahier, E. Louis, F. Baert, P. Dewint, L. Pouillon, G. Lambrecht, F. Vallée, S. Vermeire, P. Bossuyt, D. Franchimont, "Collecting New Peak and Intermediate Influximab Levels to Predict Remission in Inflammatory Bowel Diseases," in Inflammatory Bowel Diseases, 2021.